



6 Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright,  
7 no part of this document may be reproduced, stored in, or introduced into a retrieval system, or transmitted in any form  
8 or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express  
9 written permission of the authors.

10 © 2011-2018 OpenACC-Standard.org. All rights reserved.

# 11 Contents

12	<b>1. Introduction</b>	<b>7</b>
13	1.1. Scope . . . . .	7
14	1.2. Execution Model . . . . .	7
15	1.3. Memory Model . . . . .	9
16	1.4. Conventions used in this document . . . . .	11
17	1.5. Organization of this document . . . . .	11
18	1.6. References . . . . .	12
19	1.7. Changes from Version 1.0 to 2.0 . . . . .	12
20	1.8. Corrections in the August 2013 document . . . . .	13
21	1.9. Changes from Version 2.0 to 2.5 . . . . .	14
22	1.10. Changes from Version 2.5 to 2.6 . . . . .	15
23	1.11. Changes from Version 2.6 to 2.7 . . . . .	16
24	1.12. Topics Deferred For a Future Revision . . . . .	16
25	<b>2. Directives</b>	<b>19</b>
26	2.1. Directive Format . . . . .	19
27	2.2. Conditional Compilation . . . . .	20
28	2.3. Internal Control Variables . . . . .	20
29	2.3.1. Modifying and Retrieving ICV Values . . . . .	20
30	2.4. Device-Specific Clauses . . . . .	21
31	2.5. Compute Constructs . . . . .	22
32	2.5.1. Parallel Construct . . . . .	22
33	2.5.2. Kernels Construct . . . . .	23
34	2.5.3. Serial Construct . . . . .	25
35	2.5.4. if clause . . . . .	27
36	2.5.5. self clause . . . . .	27
37	2.5.6. async clause . . . . .	27
38	2.5.7. wait clause . . . . .	27
39	2.5.8. num_gangs clause . . . . .	27
40	2.5.9. num_workers clause . . . . .	27
41	2.5.10. vector_length clause . . . . .	28
42	2.5.11. private clause . . . . .	28
43	2.5.12. firstprivate clause . . . . .	28
44	2.5.13. reduction clause . . . . .	28
45	2.5.14. default clause . . . . .	29
46	2.6. Data Environment . . . . .	30
47	2.6.1. Variables with Predetermined Data Attributes . . . . .	30
48	2.6.2. Data Regions and Data Lifetimes . . . . .	31
49	2.6.3. Data Structures with Pointers . . . . .	31
50	2.6.4. Data Construct . . . . .	32

51	2.6.5. Enter Data and Exit Data Directives . . . . .	33
52	2.6.6. Reference Counters . . . . .	35
53	2.6.7. Attachment Counter . . . . .	35
54	2.7. Data Clauses . . . . .	36
55	2.7.1. Data Specification in Data Clauses . . . . .	36
56	2.7.2. Data Clause Actions . . . . .	38
57	2.7.3. deviceptr clause . . . . .	41
58	2.7.4. present clause . . . . .	41
59	2.7.5. copy clause . . . . .	42
60	2.7.6. copyin clause . . . . .	42
61	2.7.7. copyout clause . . . . .	43
62	2.7.8. create clause . . . . .	43
63	2.7.9. no_create clause . . . . .	44
64	2.7.10. delete clause . . . . .	44
65	2.7.11. attach clause . . . . .	45
66	2.7.12. detach clause . . . . .	45
67	2.8. Host_Data Construct . . . . .	45
68	2.8.1. use_device clause . . . . .	46
69	2.8.2. if clause . . . . .	46
70	2.8.3. if_present clause . . . . .	47
71	2.9. Loop Construct . . . . .	47
72	2.9.1. collapse clause . . . . .	48
73	2.9.2. gang clause . . . . .	48
74	2.9.3. worker clause . . . . .	49
75	2.9.4. vector clause . . . . .	49
76	2.9.5. seq clause . . . . .	50
77	2.9.6. auto clause . . . . .	50
78	2.9.7. tile clause . . . . .	50
79	2.9.8. device_type clause . . . . .	50
80	2.9.9. independent clause . . . . .	51
81	2.9.10. private clause . . . . .	51
82	2.9.11. reduction clause . . . . .	51
83	2.10. Cache Directive . . . . .	52
84	2.11. Combined Constructs . . . . .	53
85	2.12. Atomic Construct . . . . .	54
86	2.13. Declare Directive . . . . .	59
87	2.13.1. device_resident clause . . . . .	60
88	2.13.2. create clause . . . . .	61
89	2.13.3. link clause . . . . .	61
90	2.14. Executable Directives . . . . .	62
91	2.14.1. Init Directive . . . . .	62
92	2.14.2. Shutdown Directive . . . . .	63
93	2.14.3. Set Directive . . . . .	64
94	2.14.4. Update Directive . . . . .	65
95	2.14.5. Wait Directive . . . . .	67
96	2.14.6. Enter Data Directive . . . . .	67
97	2.14.7. Exit Data Directive . . . . .	67

98	2.15. Procedure Calls in Compute Regions . . . . .	67
99	2.15.1. Routine Directive . . . . .	68
100	2.15.2. Global Data Access . . . . .	70
101	2.16. Asynchronous Behavior . . . . .	71
102	2.16.1. async clause . . . . .	71
103	2.16.2. wait clause . . . . .	72
104	2.16.3. Wait Directive . . . . .	72
105	2.17. Fortran Optional Arguments . . . . .	73
106	<b>3. Runtime Library</b>	<b>75</b>
107	3.1. Runtime Library Definitions . . . . .	75
108	3.2. Runtime Library Routines . . . . .	76
109	3.2.1. acc_get_num_devices . . . . .	76
110	3.2.2. acc_set_device_type . . . . .	77
111	3.2.3. acc_get_device_type . . . . .	77
112	3.2.4. acc_set_device_num . . . . .	78
113	3.2.5. acc_get_device_num . . . . .	78
114	3.2.6. acc_get_property . . . . .	79
115	3.2.7. acc_init . . . . .	80
116	3.2.8. acc_shutdown . . . . .	81
117	3.2.9. acc_async_test . . . . .	81
118	3.2.10. acc_async_test_all . . . . .	82
119	3.2.11. acc_wait . . . . .	82
120	3.2.12. acc_wait_async . . . . .	83
121	3.2.13. acc_wait_all . . . . .	83
122	3.2.14. acc_wait_all_async . . . . .	84
123	3.2.15. acc_get_default_async . . . . .	84
124	3.2.16. acc_set_default_async . . . . .	85
125	3.2.17. acc_on_device . . . . .	85
126	3.2.18. acc_malloc . . . . .	86
127	3.2.19. acc_free . . . . .	86
128	3.2.20. acc_copyin . . . . .	86
129	3.2.21. acc_create . . . . .	88
130	3.2.22. acc_copyout . . . . .	89
131	3.2.23. acc_delete . . . . .	90
132	3.2.24. acc_update_device . . . . .	91
133	3.2.25. acc_update_self . . . . .	91
134	3.2.26. acc_map_data . . . . .	92
135	3.2.27. acc_unmap_data . . . . .	93
136	3.2.28. acc_deviceptr . . . . .	93
137	3.2.29. acc_hostptr . . . . .	93
138	3.2.30. acc_is_present . . . . .	94
139	3.2.31. acc_memcpy_to_device . . . . .	94
140	3.2.32. acc_memcpy_from_device . . . . .	95
141	3.2.33. acc_memcpy_device . . . . .	95
142	3.2.34. acc_attach . . . . .	96
143	3.2.35. acc_detach . . . . .	96

144	<b>4. Environment Variables</b>	<b>99</b>
145	4.1. ACC_DEVICE_TYPE . . . . .	99
146	4.2. ACC_DEVICE_NUM . . . . .	99
147	4.3. ACC_PROFLIB . . . . .	99
148	<b>5. Profiling Interface</b>	<b>101</b>
149	5.1. Events . . . . .	101
150	5.1.1. Runtime Initialization and Shutdown . . . . .	102
151	5.1.2. Device Initialization and Shutdown . . . . .	102
152	5.1.3. Enter Data and Exit Data . . . . .	103
153	5.1.4. Data Allocation . . . . .	103
154	5.1.5. Data Construct . . . . .	104
155	5.1.6. Update Directive . . . . .	104
156	5.1.7. Compute Construct . . . . .	104
157	5.1.8. Enqueue Kernel Launch . . . . .	105
158	5.1.9. Enqueue Data Update (Upload and Download) . . . . .	105
159	5.1.10. Wait . . . . .	106
160	5.2. Callbacks Signature . . . . .	106
161	5.2.1. First Argument: General Information . . . . .	107
162	5.2.2. Second Argument: Event-Specific Information . . . . .	108
163	5.2.3. Third Argument: API-Specific Information . . . . .	111
164	5.3. Loading the Library . . . . .	112
165	5.3.1. Library Registration . . . . .	113
166	5.3.2. Statically-Linked Library Initialization . . . . .	114
167	5.3.3. Runtime Dynamic Library Loading . . . . .	114
168	5.3.4. Preloading with LD_PRELOAD . . . . .	115
169	5.3.5. Application-Controlled Initialization . . . . .	116
170	5.4. Registering Event Callbacks . . . . .	116
171	5.4.1. Event Registration and Unregistration . . . . .	117
172	5.4.2. Disabling and Enabling Callbacks . . . . .	118
173	5.5. Advanced Topics . . . . .	119
174	5.5.1. Dynamic Behavior . . . . .	120
175	5.5.2. OpenACC Events During Event Processing . . . . .	121
176	5.5.3. Multiple Host Threads . . . . .	121
177	<b>6. Glossary</b>	<b>123</b>
178	<b>A. Recommendations for Implementors</b>	<b>127</b>
179	A.1. Target Devices . . . . .	127
180	A.1.1. NVIDIA GPU Targets . . . . .	127
181	A.1.2. AMD GPU Targets . . . . .	127
182	A.1.3. Multicore Host CPU Target . . . . .	128
183	A.2. API Routines for Target Platforms . . . . .	128
184	A.2.1. NVIDIA CUDA Platform . . . . .	129
185	A.2.2. OpenCL Target Platform . . . . .	130
186	A.3. Recommended Options . . . . .	131
187	A.3.1. C Pointer in Present clause . . . . .	131
188	A.3.2. Autoscopying . . . . .	131

# 1. Introduction

189

190 This document describes the compiler directives, library routines, and environment variables that  
191 collectively define the OpenACC<sup>™</sup> Application Programming Interface (OpenACC API) for writ-  
192 ing parallel programs in C, C++, and Fortran that run identified regions in parallel on multicore  
193 CPUs or attached accelerators. The method described provides a model for parallel programming  
194 that is portable across operating systems and various types of multicore CPUs and accelerators. The  
195 directives extend the ISO/ANSI standard C, C++, and Fortran base languages in a way that allows  
196 a programmer to migrate applications incrementally to parallel multicore and accelerator targets  
197 using standards-based C, C++, or Fortran.

198 The directives and programming model defined in this document allow programmers to create appli-  
199 cations capable of using accelerators without the need to explicitly manage data or program transfers  
200 between a host and accelerator or to initiate accelerator startup and shutdown. Rather, these details  
201 are implicit in the programming model and are managed by the OpenACC API-enabled compilers  
202 and runtime environments. The programming model allows the programmer to augment informa-  
203 tion available to the compilers, including specification of data local to an accelerator, guidance on  
204 mapping of loops for parallel execution, and similar performance-related details.

## 1.1. Scope

205

206 This OpenACC API document covers only user-directed parallel and accelerator programming,  
207 where the user specifies the regions of a program to be targeted for parallel execution. The remainder  
208 of the program will be executed sequentially on the host. This document does not describe features  
209 or limitations of the host programming environment as a whole; it is limited to specification of loops  
210 and regions of code to be executed in parallel on a multicore CPU or an accelerator.

211 This document does not describe automatic detection of parallel regions or automatic offloading  
212 of regions of code to an accelerator by a compiler or other tool. This document does not describe  
213 splitting loops or code regions across multiple accelerators attached to a single host. While future  
214 compilers may allow for automatic parallelization or automatic offloading, or parallelizing across  
215 multiple accelerators of the same type, or across multiple accelerators of different types, these pos-  
216 sibilities are not addressed in this document.

## 1.2. Execution Model

217

218 The execution model targeted by OpenACC API-enabled implementations is host-directed execu-  
219 tion with an attached parallel accelerator, such as a GPU, or a multicore host with a host thread that  
220 initiates parallel execution on the multiple cores, thus treating the multicore CPU itself as a device.  
221 Much of a user application executes on a host thread. Compute intensive regions are offloaded to an  
222 accelerator or executed on the multiple host cores under control of a host thread. A device, either

223 an attached accelerator or the multicore CPU, executes *parallel regions*, which typically contain  
224 work-sharing loops, *kernels regions*, which typically contain one or more loops that may be exe-  
225 cuted as kernels, or *serial regions*, which are blocks of sequential code. Even in accelerator-targeted  
226 regions, the host thread may orchestrate the execution by allocating memory on the accelerator de-  
227 vice, initiating data transfer, sending the code to the accelerator, passing arguments to the compute  
228 region, queuing the accelerator code, waiting for completion, transferring results back to the host,  
229 and deallocating memory. In most cases, the host can queue a sequence of operations to be executed  
230 on a device, one after the other.

231 Most current accelerators and many multicore CPUs support two or three levels of parallelism.  
232 Most accelerators and multicore CPUs support coarse-grain parallelism, which is fully parallel exe-  
233 cution across execution units. There may be limited support for synchronization across coarse-grain  
234 parallel operations. Many accelerators and some CPUs also support fine-grain parallelism, often  
235 implemented as multiple threads of execution within a single execution unit, which are typically  
236 rapidly switched on the execution unit to tolerate long latency memory operations. Finally, most  
237 accelerators and CPUs also support SIMD or vector operations within each execution unit. The  
238 execution model exposes these multiple levels of parallelism on a device and the programmer is  
239 required to understand the difference between, for example, a fully parallel loop and a loop that  
240 is vectorizable but requires synchronization between statements. A fully parallel loop can be pro-  
241 grammed for coarse-grain parallel execution. Loops with dependences must either be split to allow  
242 coarse-grain parallel execution, or be programmed to execute on a single execution unit using fine-  
243 grain parallelism, vector parallelism, or sequentially.

244 OpenACC exposes these three *levels of parallelism* via *gang*, *worker*, and *vector* parallelism. Gang  
245 parallelism is coarse-grain. A number of gangs will be launched on the accelerator. Worker paral-  
246 lelism is fine-grain. Each gang will have one or more workers. Vector parallelism is for SIMD or  
247 vector operations within a worker.

248 When executing a compute region on a device, one or more gangs are launched, each with one or  
249 more workers, where each worker may have vector execution capability with one or more vector  
250 lanes. The gangs start executing in *gang-redundant* mode (GR mode), meaning one vector lane of  
251 one worker in each gang executes the same code, redundantly. When the program reaches a loop  
252 or loop nest marked for gang-level work-sharing, the program starts to execute in *gang-partitioned*  
253 mode (GP mode), where the iterations of the loop or loops are partitioned across gangs for truly  
254 parallel execution, but still with only one worker per gang and one vector lane per worker active.

255 When only one worker is active, in either GR or GP mode, the program is in *worker-single* mode  
256 (WS mode). When only one vector lane is active, the program is in *vector-single* mode (VS mode).  
257 If a gang reaches a loop or loop nest marked for worker-level work-sharing, the gang transitions to  
258 *worker-partitioned* mode (WP mode), which activates all the workers of the gang. The iterations  
259 of the loop or loops are partitioned across the workers of this gang. If the same loop is marked for  
260 both gang-partitioning and worker-partitioning, then the iterations of the loop are spread across all  
261 the workers of all the gangs. If a worker reaches a loop or loop nest marked for vector-level work-  
262 sharing, the worker will transition to *vector-partitioned* mode (VP mode). Similar to WP mode, the  
263 transition to VP mode activates all the vector lanes of the worker. The iterations of the loop or loops  
264 will be partitioned across the vector lanes using vector or SIMD operations. Again, a single loop  
265 may be marked for one, two, or all three of gang, worker, and vector parallelism, and the iterations  
266 of that loop will be spread across the gangs, workers, and vector lanes as appropriate.

267 The program starts executing with a single initial host thread, identified by a program counter and



268 its stack. The initial host thread may spawn additional host threads, using OpenACC or another  
269 mechanism, such as with the OpenMP API. On a device, a single vector lane of a single worker of a  
270 single gang is called a device thread. When executing on an accelerator, a parallel execution context  
271 is created on the accelerator and may contain many such threads.

272 The user should not attempt to implement barrier synchronization, critical sections or locks across  
273 any of gang, worker, or vector parallelism. The execution model allows for an implementation that  
274 executes some gangs to completion before starting to execute other gangs. This means that trying  
275 to implement synchronization between gangs is likely to fail. In particular, a barrier across gangs  
276 cannot be implemented in a portable fashion, since all gangs may not ever be active at the same time.  
277 Similarly, the execution model allows for an implementation that executes some workers within a  
278 gang or vector lanes within a worker to completion before starting other workers or vector lanes,  
279 or for some workers or vector lanes to be suspended until other workers or vector lanes complete.  
280 This means that trying to implement synchronization across workers or vector lanes is likely to fail.  
281 In particular, implementing a barrier or critical section across workers or vector lanes using atomic  
282 operations and a busy-wait loop may never succeed, since the scheduler may suspend the worker or  
283 vector lane that owns the lock, and the worker or vector lane waiting on the lock can never complete.

284 Some devices, such as a multicore CPU, may also create and launch additional compute regions,  
285 allowing for nested parallelism. In that case, the OpenACC directives may be executed by a host  
286 thread or a device thread. This specification uses the term *local thread* or *local memory* to mean the  
287 thread that executes the directive, or the memory associated with that thread, whether that thread  
288 executes on the host or on the accelerator. The specification uses the term *local device* to mean the  
289 device on which the *local thread* is executing.

290 Most accelerators can operate asynchronously with respect to the host thread. Such devices have one  
291 or more activity queues. The host thread will enqueue operations onto the device activity queues,  
292 such as data transfers and procedure execution. After enqueueing the operation, the host thread can  
293 continue execution while the device operates independently and asynchronously. The host thread  
294 may query the device activity queue(s) and wait for all the operations in a queue to complete.  
295 Operations on a single device activity queue will complete before starting the next operation on the  
296 same queue; operations on different activity queues may be active simultaneously and may complete  
297 in any order.

## 298 1.3. Memory Model

299 The most significant difference between a host-only program and a host+accelerator program is that  
300 the memory on an accelerator may be discrete from host memory. This is the case with most current  
301 GPUs, for example. In this case, the host thread may not be able to read or write device memory  
302 directly because it is not mapped into the host thread's virtual memory space. All data movement  
303 between host memory and accelerator memory must be performed by the host thread through system  
304 calls that explicitly move data between the separate memories, typically using direct memory access  
305 (DMA) transfers. Similarly, it is not valid to assume the accelerator can read or write host memory,  
306 though this is supported by some accelerators, often with significant performance penalty.

307 The concept of discrete host and accelerator memories is very apparent in low-level accelerator  
308 programming languages such as CUDA or OpenCL, in which data movement between the memories  
309 can dominate user code. In the OpenACC model, data movement between the memories can be  
310 implicit and managed by the compiler, based on directives from the programmer. However, the

311 programmer must be aware of the potentially discrete memories for many reasons, including but  
312 not limited to:

- 313 • Memory bandwidth between host memory and accelerator memory determines the level of  
314 compute intensity required to effectively accelerate a given region of code.
- 315 • The user should be aware that a discrete device memory is usually significantly smaller than  
316 the host memory, prohibiting offloading regions of code that operate on very large amounts  
317 of data.
- 318 • Host addresses stored to pointers on the host may only be valid on the host; addresses stored  
319 to pointers in accelerator memory may only be valid on that device. Explicitly transferring  
320 pointer values between host and accelerator memory is not advised. Dereferencing host point-  
321 ers on an accelerator or dereferencing accelerator pointers on the host is likely to be invalid  
322 on such targets.

323 OpenACC exposes the discrete memories through the use of a device data environment. Device data  
324 has an explicit lifetime, from when it is allocated or created until it is deleted. If a device shares  
325 memory with the local thread, its device data environment will be shared with the local thread. In  
326 that case, the implementation need not create new copies of the data for the device and no data  
327 movement need be done. If a device has a discrete memory and shares no memory with the local  
328 thread, the implementation will allocate space in device memory and copy data between the local  
329 memory and device memory, as appropriate. The local thread may share some memory with a  
330 device and also have some memory that is not shared with that device. In that case, data in shared  
331 memory may be accessed by both the local thread and the device. Data not in shared memory will  
332 be copied to device memory as necessary.

333 Some accelerators (such as current GPUs) implement a weak memory model. In particular, they do  
334 not support memory coherence between operations executed by different threads; even on the same  
335 execution unit, memory coherence is only guaranteed when the memory operations are separated  
336 by an explicit memory fence. Otherwise, if one thread updates a memory location and another reads  
337 the same location, or two threads store a value to the same location, the hardware may not guarantee  
338 the same result for each execution. While a compiler can detect some potential errors of this nature,  
339 it is nonetheless possible to write a compute region that produces inconsistent numerical results.

340 Similarly, some accelerators implement a weak memory model for memory shared between the  
341 host and the accelerator, or memory shared between multiple accelerators. Programmers need to  
342 be very careful that the program uses appropriate synchronization to ensure that an assignment or  
343 modification by a thread on any device to data in shared memory is complete and available before  
344 that data is used by another thread on the same or another device.

345 Some current accelerators have a software-managed cache, some have hardware managed caches,  
346 and most have hardware caches that can be used only in certain situations and are limited to read-  
347 only data. In low-level programming models such as CUDA or OpenCL languages, it is up to the  
348 programmer to manage these caches. In the OpenACC model, these caches are managed by the  
349 compiler with hints from the programmer in the form of directives.

## 1.4. Conventions used in this document

Some terms are used in this specification that conflict with their usage as defined in the base languages. When there is potential confusion, the term will appear in the Glossary.

Keywords and punctuation that are part of the actual specification will appear in typewriter font:

```
#pragma acc
```

Italic font is used where a keyword or other name must be used:

```
#pragma acc directive-name
```

For C and C++, *new-line* means the newline character at the end of a line:

```
#pragma acc directive-name new-line
```

Optional syntax is enclosed in square brackets; an option that may be repeated more than once is followed by ellipses:

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

In this spec, a *var* (in italics) is one of the following:

- a variable name (a scalar, array, or composite variable name);
- a subarray specification with subscript ranges;
- an array element;
- a member of a composite variable;
- a common block name between slashes.

Not all options are allowed in all clauses; the allowable options are clarified for each use of the term *var*.

To simplify the specification and convey appropriate constraint information, a *pqr-list* is a comma-separated list of *pqr* items. For example, an *int-expr-list* is a comma-separated list of one or more integer expressions, and a *var-list* is a comma-separated list of one or more *vars*. The one exception is *clause-list*, which is a list of one or more clauses optionally separated by commas.

```
#pragma acc directive-name [clause-list] new-line
```

## 1.5. Organization of this document

The rest of this document is organized as follows:

Chapter 2 Directives, describes the C, C++, and Fortran directives used to delineate accelerator regions and augment information available to the compiler for scheduling of loops and classification of data.

375 Chapter 3 Runtime Library, defines user-callable functions and library routines to query the accel-  
376 erator features and control behavior of accelerator-enabled programs at runtime.

377 Chapter 4 Environment Variables, defines user-settable environment variables used to control be-  
378 havior of accelerator-enabled programs at execution.

379 Chapter 5 Profiling Interface, describes the OpenACC interface for tools that can be used for profile  
380 and trace data collection.

381 Chapter 6 Glossary, defines common terms used in this document.

382 Appendix A Recommendations for Implementors, gives advice to implementers to support more  
383 portability across implementations and interoperability with other accelerator APIs.

## 384 1.6. References

- 385 • *American National Standard Programming Language C*, ANSI X3.159-1989 (ANSI C).
- 386 • ISO/IEC 9899:1999, *Information Technology – Programming Languages – C (C99)*.
- 387 • ISO/IEC 14882:1998, *Information Technology – Programming Languages – C++*.
- 388 • ISO/IEC 1539-1:2004, *Information Technology – Programming Languages – Fortran – Part*  
389 *1: Base Language*, (Fortran 2003).
- 390 • *OpenMP Application Program Interface*, version 4.0, July 2013
- 391 • *PGI Accelerator Programming Model for Fortran & C*, version 1.3, November 2011
- 392 • *NVIDIA CUDA<sup>TM</sup> C Programming Guide*, version 7.0, March 2015.
- 393 • *The OpenCL Specification*, version 202, Khronos OpenCL Working Group, October 2014.

## 394 1.7. Changes from Version 1.0 to 2.0

- 395 • `_OPENACC` value updated to `201306`
- 396 • `default (none)` clause on `parallel` and `kernels` directives
- 397 • the implicit data attribute for scalars in `parallel` constructs has changed
- 398 • the implicit data attribute for scalars in loops with `loop` directives with the independent  
399 attribute has been clarified
- 400 • `acc_async_sync` and `acc_async_noval` values for the `async` clause
- 401 • Clarified the behavior of the `reduction` clause on a `gang` loop
- 402 • Clarified allowable loop nesting (`gang` may not appear inside `worker`, which may not ap-  
403 pear within `vector`)
- 404 • `wait` clause on `parallel`, `kernels` and `update` directives
- 405 • `async` clause on the `wait` directive
- 406 • `enter data` and `exit data` directives

- 407 • Fortran *common block* names may now be specified in many data clauses
- 408 • **link** clause for the **declare** directive
- 409 • the behavior of the **declare** directive for global data
- 410 • the behavior of a data clause with a C or C++ pointer variable has been clarified
- 411 • predefined data attributes
- 412 • support for multidimensional dynamic C/C++ arrays
- 413 • **tile** and **auto** loop clauses
- 414 • **update self** introduced as a preferred synonym for **update host**
- 415 • **routine** directive and support for separate compilation
- 416 • **device\_type** clause and support for multiple device types
- 417 • nested parallelism using parallel or kernels region containing another parallel or kernels re-
- 418 gion
- 419 • **atomic** constructs
- 420 • new concepts: gang-redundant, gang-partitioned; worker-single, worker-partitioned; vector-
- 421 single, vector-partitioned; thread
- 422 • new API routines:
  - 423 – **acc\_wait**, **acc\_wait\_all** instead of **acc\_async\_wait** and **acc\_async\_wait\_all**
  - 424 – **acc\_wait\_async**
  - 425 – **acc\_copyin**, **acc\_present\_or\_copyin**
  - 426 – **acc\_create**, **acc\_present\_or\_create**
  - 427 – **acc\_copyout**, **acc\_delete**
  - 428 – **acc\_map\_data**, **acc\_unmap\_data**
  - 429 – **acc\_deviceptr**, **acc\_hostptr**
  - 430 – **acc\_is\_present**
  - 431 – **acc\_memcpy\_to\_device**, **acc\_memcpy\_from\_device**
  - 432 – **acc\_update\_device**, **acc\_update\_self**
- 433 • defined behavior with multiple host threads, such as with OpenMP
- 434 • recommendations for specific implementations
- 435 • clarified that no arguments are allowed on the **vector** clause in a parallel region

## 436 1.8. Corrections in the August 2013 document

- 437 • corrected the **atomic capture** syntax for C/C++
- 438 • fixed the name of the **acc\_wait** and **acc\_wait\_all** procedures
- 439 • fixed description of the **acc\_hostptr** procedure

## 1.9. Changes from Version 2.0 to 2.5

- The `_OPENACC` value was updated to `201510`; see Section 2.2 Conditional Compilation.
- The `num_gangs`, `num_workers`, and `vector_length` clauses are now allowed on the `kernels` construct; see Section 2.5.2 Kernels Construct.
- Reduction on C++ class members, array elements, and struct elements are explicitly disallowed; see Section 2.5.13 reduction clause.
- Reference counting is now used to manage the correspondence and lifetime of device data; see Section 2.6.6 Reference Counters.
- The behavior of the `exit data` directive has changed to decrement the dynamic reference counter. A new optional `finalize` clause was added to set the dynamic reference counter to zero. See Section 2.6.5 Enter Data and Exit Data Directives.
- The `copy`, `copyin`, `copyout`, and `create` data clauses were changed to behave like `present_or_copy`, etc. The `present_or_copy`, `pcopy`, `present_or_copyin`, `pcopyin`, `present_or_copyout`, `pcopyout`, `present_or_create`, and `pcreate` data clauses are no longer needed, though will be accepted for compatibility; see Section 2.7 Data Clauses.
- Reductions on orphaned gang loops are explicitly disallowed; see Section 2.9 Loop Construct.
- The description of the `loop auto` clause has changed; see Section 2.9.6 auto clause.
- Text was added to the `private` clause on a `loop` construct to clarify that a copy is made for each gang or worker or vector lane, not each thread; see Section 2.9.10 private clause.
- The description of the `reduction` clause on a `loop` construct was corrected; see Section 2.9.11 reduction clause.
- A restriction was added to the `cache` clause that all references to that variable must lie within the region being cached; see Section 2.10 Cache Directive.
- Text was added to the `private` and `reduction` clauses on a combined construct to clarify that they act like `private` and `reduction` on the `loop`, not `private` and `reduction` on the `parallel` or `reduction` on the `kernels`; see Section 2.11 Combined Constructs.
- The `declare create` directive with a Fortran `allocatable` has new behavior; see Section 2.13.2 create clause.
- New `init`, `shutdown`, `set` directives were added; see Section 2.14.1 Init Directive, 2.14.2 Shutdown Directive, and 2.14.3 Set Directive.
- A new `if_present` clause was added to the `update` directive, which changes the behavior when data is not present from a runtime error to a no-op; see Section 2.14.4 Update Directive.
- The `routine bind` clause definition changed; see Section 2.15.1 Routine Directive.
- An `acc routine` without `gang/worker/vector/seq` is now defined as an error; see Section 2.15.1 Routine Directive.
- A new `default (present)` clause was added for compute constructs; see Section 2.5.14 default clause.
- The Fortran header file `openacc_lib.h` is no longer supported; the Fortran module `openacc`

- 479 should be used instead; see Section 3.1 Runtime Library Definitions.
- 480 • New API routines were added to get and set the default async queue value; see Section 3.2.15
  - 481 `acc_get_default_async` and 3.2.16 `acc_set_default_async`.
  - 482 • The `acc_copyin`, `acc_create`, `acc_copyout`, and `acc_delete` API routines were
  - 483 changed to behave like `acc_present_or_copyin`, etc. The `acc_present_or_` names
  - 484 are no longer needed, though will be supported for compatibility. See Sections 3.2.20 and fol-
  - 485 lowing.
  - 486 • Asynchronous versions of the data API routines were added; see Sections 3.2.20 and follow-
  - 487 ing.
  - 488 • A new API routine added, `acc_memcpy_device`, to copy from one device address to
  - 489 another device address; see Section 3.2.31 `acc_memcpy_to_device`.
  - 490 • A new OpenACC interface for profile and trace tools was added; see Chapter 5 Profiling Interface.

## 491 1.10. Changes from Version 2.5 to 2.6

- 492 • The `_OPENACC` value was updated to `201711`.
- 493 • A new `serial` compute construct was added. See Section 2.5.3 Serial Construct.
- 494 • A new runtime API query routine was added. `acc_get_property` may be called from
- 495 the host and returns properties about any device. See Section 3.2.6.
- 496 • The text has clarified that if a variable is in a reduction which spans two or more nested loops,
- 497 each `loop` directive on any of those loops must have a `reduction` clause that contains the
- 498 variable; see Section 2.9.11 reduction clause.
- 499 • An optional `if` or `if_present` clause is now allowed on the `host_data` construct. See
- 500 Section 2.8 Host Data Construct.
- 501 • A new `no_create` data clause is now allowed on compute and `data` constructs. See Sec-
- 502 tion 2.7.9 `no_create` clause.
- 503 • The behavior of Fortran optional arguments in data clauses and in routine calls has been
- 504 specified; see Section 2.17 Fortran Optional Arguments.
- 505 • The descriptions of some of the Fortran versions of the runtime library routines were simpli-
- 506 fied; see Section 3.2 Runtime Library Routines.
- 507 • To allow for manual deep copy of data structures with pointers, new `attach` and `detach` be-
- 508 havior was added to the data clauses, new `attach` and `detach` clauses were added, and
- 509 matching `acc_attach` and `acc_detach` runtime API routines were added; see Sections
- 510 2.6.3, 2.7.11-2.7.12 and 3.2.34-3.2.35.
- 511 • The Intel Coprocessor Offload Interface target and API routine sections were removed from
- 512 the Section A Recommendations for Implementors, since Intel no longer produces this prod-
- 513 uct.

## 1.11. Changes from Version 2.6 to 2.7

- The `_OPENACC` value was updated to **201811**.
- The specification allows for hosts that share some memory with the device but not all memory. The wording in the text now discusses whether local thread data is in shared memory (memory shared between the local thread and the device) or discrete memory (local thread memory that is not shared with the device), instead of shared-memory devices and non-shared memory devices. See Sections 1.3 Memory Model and 2.6 Data Environment.
- The text was clarified to allow an implementation that treats a multicore CPU as a device, either an additional device or the only device.
- The `readonly` modifier was added to the `copyin` data clause and `cache` directive. See Sections 2.7.6 and 2.10.
- The term *local device* was defined; see Section 1.2 Execution Model and the Glossary.
- The term *var* is used more consistently throughout the specification to mean a variable name, array name, subarray specification, array element, composite variable member, or Fortran common block name between slashes. Some uses of *var* allow only a subset of these options, and those limitations are given in those cases.
- The `self` clause was added to the compute constructs; see Section 2.5.5 self clause.
- The appearance of a `reduction` clause on a compute construct implies a `copy` clause for each reduction variable; see Sections 2.5.13 reduction clause and 2.11 Combined Constructs.
- The `default (none)` and `default (present)` clauses were added to the `data` construct; see Section 2.6.4 Data Construct.
- Data is defined to be *present* based on the values of the structured and dynamic reference counters; see Section 2.6.6 Reference Counters and the Glossary.
- The interaction of the `acc_map_data` and `acc_unmap_data` runtime API calls on the present counters is defined; see Section 2.7.2, 3.2.26, and 3.2.27.
- A restriction clarifying that a `host_data` construct must have at least one `use_device` clause was added.
- Arrays, subarrays and composite variables are now allowed in `reduction` clauses; see Sections 2.9.11 reduction clause and 2.5.13 reduction clause.
- Changed behavior of ICVs to support nested compute regions and host as a device semantics. See Section 2.3.

## 1.12. Topics Deferred For a Future Revision

The following topics are under discussion for a future revision. Some of these are known to be important, while others will depend on feedback from users. Readers who have feedback or want to participate may post a message at the forum at [www.openacc.org](http://www.openacc.org), or may send email to [technical@openacc.org](mailto:technical@openacc.org) or [feedback@openacc.org](mailto:feedback@openacc.org). No promises are made or implied that all these items will be available in the next revision.



- 551 • Support for attaching C/C++ pointers that point to an address past the end of a memory region.
- 552 • Full support for C and C++ structs and struct members, including pointer members.
- 553 • Full support for Fortran derived types and derived type members, including allocatable and  
554 pointer members.
- 555 • Fully defined interaction with multiple host threads.
- 556 • Optionally removing the synchronization or barrier at the end of vector and worker loops.
- 557 • Allowing an **if** clause after a **device\_type** clause.
- 558 • A **shared** clause (or something similar) for the loop directive.
- 559 • Better support for multiple devices from a single thread, whether of the same type or of  
560 different types.



## 2. Directives

This chapter describes the syntax and behavior of the OpenACC directives. In C and C++, OpenACC directives are specified using the **#pragma** mechanism provided by the language. In Fortran, OpenACC directives are specified using special comments that are identified by a unique sentinel. Compilers will typically ignore OpenACC directives if support is disabled or not provided.

### 2.1. Directive Format

In C and C++, OpenACC directives are specified with the **#pragma** mechanism. The syntax of an OpenACC directive is:

```
#pragma acc directive-name [clause-list] new-line
```

Each directive starts with **#pragma acc**. The remainder of the directive follows the C and C++ conventions for pragmas. White space may be used before and after the **#**; white space may be required to separate words in a directive. Preprocessing tokens following the **#pragma acc** are subject to macro replacement. Directives are case-sensitive.

In Fortran, OpenACC directives are specified in free-form source files as

```
!$acc directive-name [clause-list]
```

The comment prefix (**!**) may appear in any column, but may only be preceded by white space (spaces and tabs). The sentinel (**!\$acc**) must appear as a single word, with no intervening white space. Line length, white space, and continuation rules apply to the directive line. Initial directive lines must have white space after the sentinel. Continued directive lines must have an ampersand (**&**) as the last nonblank character on the line, prior to any comment placed in the directive. Continuation directive lines must begin with the sentinel (possibly preceded by white space) and may have an ampersand as the first non-white space character after the sentinel. Comments may appear on the same line as a directive, starting with an exclamation point and extending to the end of the line. If the first nonblank character after the sentinel is an exclamation point, the line is ignored.

In Fortran fixed-form source files, OpenACC directives are specified as one of

```
!$acc directive-name [clause-list]  
c$acc directive-name [clause-list]  
*$acc directive-name [clause-list]
```

The sentinel (**!\$acc**, **c\$acc**, or **\*\$acc**) must occupy columns 1-5. Fixed form line length, white space, continuation, and column rules apply to the directive line. Initial directive lines must have

586 a space or zero in column 6, and continuation directive lines must have a character other than a  
587 space or zero in column 6. Comments may appear on the same line as a directive, starting with an  
588 exclamation point on or after column 7 and continuing to the end of the line.

589 In Fortran, directives are case-insensitive. Directives cannot be embedded within continued state-  
590 ments, and statements must not be embedded within continued directives. In this document, free  
591 form is used for all Fortran OpenACC directive examples.

592 Only one *directive-name* can be specified per directive, except that a combined directive name is  
593 considered a single *directive-name*. The order in which clauses appear is not significant unless  
594 otherwise specified. Clauses may be repeated unless otherwise specified. Some clauses have an  
595 argument that can contain a list.

## 596 2.2. Conditional Compilation

597 The `_OPENACC` macro name is defined to have a value `yyyymm` where `yyyy` is the year and `mm` is  
598 the month designation of the version of the OpenACC directives supported by the implementation.  
599 This macro must be defined by a compiler only when OpenACC directives are enabled. The version  
600 described here is 201811.

## 601 2.3. Internal Control Variables

602 An OpenACC implementation acts as if there are internal control variables (ICVs) that control the  
603 behavior of the program. These ICVs are initialized by the implementation, and may be given  
604 values through environment variables and through calls to OpenACC API routines. The program  
605 can retrieve values through calls to OpenACC API routines.

606 The ICVs are:

- 607 • *acc-current-device-type-var* - controls which type of device is used.
- 608 • *acc-current-device-num-var* - controls which device of the selected type is used.
- 609 • *acc-default-async-var* - controls which asynchronous queue is used when none is specified in  
610 an `async` clause.

### 611 2.3.1. Modifying and Retrieving ICV Values

612 The following table shows environment variables or procedures to modify the values of the internal  
613 control variables, and procedures to retrieve the values:

ICV	Ways to modify values	Way to retrieve value
<i>acc-current-device-type-var</i>	<b>acc_set_device_type</b> <b>set device_type</b> <b>ACC_DEVICE_TYPE</b>	<b>acc_get_device_type</b>
614 <i>acc-current-device-num-var</i>	<b>acc_set_device_num</b> <b>set device_num</b> <b>ACC_DEVICE_NUM</b>	<b>acc_get_device_num</b>
<i>acc-default-async-var</i>	<b>acc_set_default_async</b> <b>set default_async</b>	<b>acc_get_default_async</b>

615 The initial values are implementation-defined. After initial values are assigned, but before any  
616 OpenACC construct or API routine is executed, the values of any environment variables that were  
617 set by the user are read and the associated ICVs are modified accordingly. There is one copy of  
618 each ICV for each host thread that is not generated by a compute construct. For threads that are  
619 generated by a compute construct the initial value for each ICV is inherited from the local thread.  
620 The behavior for each ICV is as if there is a copy for each thread. If an ICV is modified, then a  
621 unique copy of that ICV must be created for the modifying thread.

## 622 2.4. Device-Specific Clauses

623 OpenACC directives can specify different clauses or clause arguments for different devices using the  
624 **device\_type** clause. The argument to the **device\_type** is a comma-separated list of one or  
625 more device architecture name identifiers, or an asterisk. A single directive may have one or several  
626 **device\_type** clauses. Clauses on a directive with no **device\_type** apply to all device types.  
627 Clauses that follow a **device\_type** up to the end of the directive or up to the next **device\_type**  
628 are associated with this **device\_type**. Clauses associated with a **device\_type** apply only  
629 when compiling for the device type named. Clauses associated with a **device\_type** that has  
630 an asterisk argument apply to any device type that was not named in any **device\_type** on that  
631 directive. The **device\_type** clauses may appear in any order. For each directive, only certain  
632 clauses may follow a **device\_type**.

633 Clauses that precede any **device\_type** are *default clauses*. Clauses that follow a **device\_type**  
634 are *device-specific clauses*. A clause may appear both as a default clause and as a device-specific  
635 clause. In that case, the value in the device-specific clause is used when compiling for that device  
636 type.

637 The supported device types are implementation-defined. Depending on the implementation and the  
638 compiling environment, an implementation may support only a single device type, or may support  
639 multiple device types but only one at a time, or many support multiple device types in a single  
640 compilation.

641 A device architecture name may be generic, such as a vendor, or more specific, such as a partic-  
642 ular generation of device; see Appendix A Recommendations for Implementors for recommended  
643 names. When compiling for a particular device, the implementation will use the clauses associated  
644 with the **device\_type** clause that specifies the most specific architecture name that applies for  
645 this device; clauses associated with any other **device\_type** clause are ignored. In this context,  
646 the asterisk is the least specific architecture name.

647 **Syntax** The syntax of the `device_type` clause is

```
device_type( * )
device_type( device-type-list )
```

648 The `device_type` clause may be abbreviated to `dtype`.

## 649 2.5. Compute Constructs

### 650 2.5.1. Parallel Construct

651 **Summary** This fundamental construct starts parallel execution on the current device.

652 **Syntax** In C and C++, the syntax of the OpenACC `parallel` construct is

```
#pragma acc parallel [clause-list] new-line
    structured block
```

653 and in Fortran, the syntax is

```
!$acc parallel [clause-list]
    structured block
!$acc end parallel
```

654 where *clause* is one of the following:

```
async [( int-expr )]
wait [( int-expr-list )]
num_gangs( int-expr )
num_workers( int-expr )
vector_length( int-expr )
device_type( device-type-list )
if( condition )
self [( condition )]
reduction( operator:var-list )
copy( var-list )
copyin( [readonly:]var-list )
copyout( var-list )
create( var-list )
no_create( var-list )
present( var-list )
deviceptr( var-list )
attach( var-list )
private( var-list )
firstprivate( var-list )
default( none | present )
```

655 **Description** When the program encounters an accelerator **parallel** construct, one or more  
 656 gangs of workers are created to execute the accelerator parallel region. The number of gangs, and  
 657 the number of workers in each gang and the number of vector lanes per worker remain constant for  
 658 the duration of that parallel region. Each gang begins executing the code in the structured block  
 659 in gang-redundant mode. This means that code within the parallel region, but outside of a loop  
 660 construct with gang-level worksharing, will be executed redundantly by all gangs.

661 One worker in each gang begins executing the code in the structured block of the construct. Note:  
 662 Unless there is a **loop** construct within the parallel region, all gangs will execute all the code within  
 663 the region redundantly.

664 If the **async** clause does not appear, there is an implicit barrier at the end of the accelerator parallel  
 665 region, and the execution of the local thread will not proceed until all gangs have reached the end  
 666 of the parallel region.

667 If there is no **default (none)** clause on the construct, the compiler will implicitly determine data  
 668 attributes for variables that are referenced in the compute construct that do not have predetermined  
 669 data attributes and do not appear in a data clause on the compute construct, a lexically containing  
 670 **data** construct, or a visible **declare** directive. If there is no **default (present)** clause  
 671 on the construct, an array or composite variable referenced in the **parallel** construct that does  
 672 not appear in a data clause for the construct or any enclosing **data** construct will be treated as if  
 673 it appeared in a **copy** clause for the **parallel** construct. If there is a **default (present)**  
 674 clause on the construct, the compiler will implicitly treat all arrays and composite variables without  
 675 predetermined data attributes as if they appeared in a **present** clause. A scalar variable referenced  
 676 in the **parallel** construct that does not appear in a data clause for the construct or any enclosing  
 677 **data** construct will be treated as if it appeared in a **firstprivate** clause.

## 678 Restrictions

- 679 • A program may not branch into or out of an OpenACC **parallel** construct.
- 680 • A program must not depend on the order of evaluation of the clauses, or on any side effects  
 681 of the evaluations.
- 682 • Only the **async**, **wait**, **num\_gangs**, **num\_workers**, and **vector\_length** clauses  
 683 may follow a **device\_type** clause.
- 684 • At most one **if** clause may appear. In Fortran, the condition must evaluate to a scalar logical  
 685 value; in C or C++, the condition must evaluate to a scalar integer value.
- 686 • At most one **default** clause may appear, and it must have a value of either **none** or  
 687 **present**.

688 The **copy**, **copyin**, **copyout**, **create**, **no\_create**, **present**, **deviceptr**, and **attach**  
 689 data clauses are described in Section 2.7 Data Clauses. The **private** and **firstprivate**  
 690 clauses are described in Sections 2.5.11 and Sections 2.5.12. The **device\_type** clause is de-  
 691 scribed in Section 2.4 Device-Specific Clauses.

## 692 2.5.2. Kernels Construct

693 **Summary** This construct defines a region of the program that is to be compiled into a sequence  
 694 of kernels for execution on the current device.

695 **Syntax** In C and C++, the syntax of the OpenACC **kernels** construct is

```
#pragma acc kernels [clause-list] new-line
    structured block
```

696 and in Fortran, the syntax is

```
!$acc kernels [clause-list]
    structured block
!$acc end kernels
```

697 where *clause* is one of the following:

```
async [ ( int-expr ) ]
wait [ ( int-expr-list ) ]
num_gangs ( int-expr )
num_workers ( int-expr )
vector_length ( int-expr )
device_type ( device-type-list )
if ( condition )
self [ ( condition ) ]
copy ( var-list )
copyin ( [readonly:]var-list )
copyout ( var-list )
create ( var-list )
no_create ( var-list )
present ( var-list )
deviceptr ( var-list )
attach ( var-list )
default ( none | present )
```

698 **Description** The compiler will split the code in the kernels region into a sequence of acceler-  
 699 ator kernels. Typically, each loop nest will be a distinct kernel. When the program encounters a  
 700 **kernels** construct, it will launch the sequence of kernels in order on the device. The number and  
 701 configuration of gangs of workers and vector length may be different for each kernel.

702 If the **async** clause does not appear, there is an implicit barrier at the end of the kernels region, and  
 703 the local thread execution will not proceed until all kernels have completed execution.

704 If there is no **default (none)** clause on the construct, the compiler will implicitly determine data  
 705 attributes for variables that are referenced in the compute construct that do not have predetermined  
 706 data attributes and do not appear in a data clause on the compute construct, a lexically containing  
 707 **data** construct, or a visible **declare** directive. If there is no **default (present)** clause  
 708 on the construct, an array or composite variable referenced in the **kernels** construct that does  
 709 not appear in a data clause for the construct or any enclosing **data** construct will be treated as  
 710 if it appeared in a **copy** clause for the **kernels** construct. If there is a **default (present)**  
 711 clause on the construct, the compiler will implicitly treat all arrays and composite variables without  
 712 predetermined data attributes as if they appeared in a **present** clause. A scalar variable referenced



713 in the **kernels** construct that does not appear in a data clause for the construct or any enclosing  
 714 **data** construct will be treated as if it appeared in a **copy** clause.

## 715 Restrictions

- 716 • A program may not branch into or out of an OpenACC **kernels** construct.
- 717 • A program must not depend on the order of evaluation of the clauses, or on any side effects  
 718 of the evaluations.
- 719 • Only the **async**, **wait**, **num\_gangs**, **num\_workers**, and **vector\_length** clauses  
 720 may follow a **device\_type** clause.
- 721 • At most one **if** clause may appear. In Fortran, the condition must evaluate to a scalar logical  
 722 value; in C or C++, the condition must evaluate to a scalar integer value.
- 723 • At most one **default** clause may appear, and it must have a value of either **none** or  
 724 **present**.

725 The **copy**, **copyin**, **copyout**, **create**, **no\_create**, **present**, **deviceptr**, and **attach**  
 726 data clauses are described in Section 2.7 Data Clauses. The **device\_type** clause is described in  
 727 Section 2.4 Device-Specific Clauses.

## 728 2.5.3. Serial Construct

729 **Summary** This construct defines a region of the program that is to be executed sequentially on  
 730 the current device.

731 **Syntax** In C and C++, the syntax of the OpenACC **serial** construct is

```
#pragma acc serial [clause-list] new-line
      structured block
```

732 and in Fortran, the syntax is

```
!$acc serial [clause-list]
      structured block
!$acc end serial
```

733 where *clause* is one of the following:

```
async [( int-expr )]
wait [( int-expr-list )]
device_type( device-type-list )
if( condition )
self [( condition )]
reduction( operator:var-list )
copy( var-list )
```

```

copyin( [readonly:]var-list )
copyout( var-list )
create( var-list )
no_create( var-list )
present( var-list )
deviceptr( var-list )
private( var-list )
firstprivate( var-list )
attach( var-list )
default( none | present )

```

734 **Description** When the program encounters an accelerator **serial** construct, one gang of one  
 735 worker with a vector length of one is created to execute the accelerator serial region sequentially.  
 736 The single gang begins executing the code in the structured block in gang-redundant mode, even  
 737 though there is a single gang. The **serial** construct executes as if it were a **parallel** construct  
 738 with clauses **num\_gangs(1) num\_workers(1) vector\_length(1)**.

739 If the **async** clause does not appear, there is an implicit barrier at the end of the accelerator serial  
 740 region, and the execution of the local thread will not proceed until the gang has reached the end of  
 741 the serial region.

742 If there is no **default(none)** clause on the construct, the compiler will implicitly determine data  
 743 attributes for variables that are referenced in the compute construct that do not have predetermined  
 744 data attributes and do not appear in a data clause on the compute construct, a lexically containing  
 745 **data** construct, or a visible **declare** directive. If there is no **default(present)** clause  
 746 on the construct, an array or composite variable referenced in the **serial** construct that does  
 747 not appear in a data clause for the construct or any enclosing **data** construct will be treated as  
 748 if it appeared in a **copy** clause for the **serial** construct. If there is a **default(present)**  
 749 clause on the construct, the compiler will implicitly treat all arrays and composite variables without  
 750 predetermined data attributes as if they appeared in a **present** clause. A scalar variable referenced  
 751 in the **serial** construct that does not appear in a data clause for the construct or any enclosing  
 752 **data** construct will be treated as if it appeared in a **firstprivate** clause.

### 753 Restrictions

- 754 • A program may not branch into or out of an OpenACC **serial** construct.
- 755 • A program must not depend on the order of evaluation of the clauses, or on any side effects  
756 of the evaluations.
- 757 • Only the **async** and **wait** clauses may follow a **device\_type** clause.
- 758 • At most one **if** clause may appear. In Fortran, the condition must evaluate to a scalar logical  
759 value; in C or C++, the condition must evaluate to a scalar integer value.
- 760 • At most one **default** clause may appear, and it must have a value of either **none** or  
761 **present**.

762 The **copy**, **copyin**, **copyout**, **create**, **no\_create**, **present**, **deviceptr**, and **attach**  
 763 data clauses are described in Section 2.7 Data Clauses. The **private** and **firstprivate**  
 764 clauses are described in Sections 2.5.11 and Sections 2.5.12. The **device\_type** clause is de-

765 scribed in Section 2.4 Device-Specific Clauses.

#### 766 **2.5.4. if clause**

767 The **if** clause is optional.

768 When the *condition* in the **if** clause evaluates to nonzero in C or C++, or **.true.** in Fortran, the  
769 region will execute on the current device. When the *condition* in the **if** clause evaluates to zero in  
770 C or C++, or **.false.** in Fortran, the local thread will execute the region.

#### 771 **2.5.5. self clause**

772 The **self** clause is optional.

773 The **self** clause may have a single *condition-argument*. If the *condition-argument* is not present  
774 it is assumed to be nonzero in C or C++, or **.true.** in Fortran. When both an **if** clause and a  
775 **self** clause appear and the *condition* in the **if** clause evaluates to 0 in C or C++ or **.false.** in  
776 Fortran, the **self** clause has no effect.

777 When the *condition* evaluates to nonzero in C or C++, or **.true.** in Fortran, the region will execute  
778 on the local device. When the *condition* in the **self** clause evaluates to zero in C or C++, or  
779 **.false.** in Fortran, the region will execute on the current device.

#### 780 **2.5.6. async clause**

781 The **async** clause is optional; see Section 2.16 Asynchronous Behavior for more information.

#### 782 **2.5.7. wait clause**

783 The **wait** clause is optional; see Section 2.16 Asynchronous Behavior for more information.

#### 784 **2.5.8. num\_gangs clause**

785 The **num\_gangs** clause is allowed on the **parallel** and **kernels** constructs. The value of  
786 the integer expression defines the number of parallel gangs that will execute the parallel region,  
787 or that will execute each kernel created for the kernels region. If the clause is not specified, an  
788 implementation-defined default will be used; the default may depend on the code within the con-  
789 struct. The implementation may use a lower value than specified based on limitations imposed by  
790 the target architecture.

#### 791 **2.5.9. num\_workers clause**

792 The **num\_workers** clause is allowed on the **parallel** and **kernels** constructs. The value  
793 of the integer expression defines the number of workers within each gang that will be active af-  
794 ter a gang transitions from worker-single mode to worker-partitioned mode. If the clause is not  
795 specified, an implementation-defined default will be used; the default value may be 1, and may be

796 different for each **parallel** construct or for each kernel created for a **kernels** construct. The  
797 implementation may use a different value than specified based on limitations imposed by the target  
798 architecture.

### 799 2.5.10. **vector\_length** clause

800 The **vector\_length** clause is allowed on the **parallel** and **kernels** constructs. The value  
801 of the integer expression defines the number of vector lanes that will be active after a worker transi-  
802 tions from vector-single mode to vector-partitioned mode. This clause determines the vector length  
803 to use for vector or SIMD operations. If the clause is not specified, an implementation-defined de-  
804 fault will be used. This vector length will be used for loop constructs annotated with the **vector**  
805 clause, as well as loops automatically vectorized by the compiler. The implementation may use a  
806 different value than specified based on limitations imposed by the target architecture.

### 807 2.5.11. **private** clause

808 The **private** clause is allowed on the **parallel** and **serial** constructs; it declares that a copy  
809 of each item on the list will be created for each gang.

#### 810 Restrictions

- 811 • See Section 2.17 Fortran Optional Arguments for discussion of Fortran optional arguments in  
812 **private** clauses.

### 813 2.5.12. **firstprivate** clause

814 The **firstprivate** clause is allowed on the **parallel** and **serial** constructs; it declares that  
815 a copy of each item on the list will be created for each gang, and that the copy will be initialized with  
816 the value of that item on the local thread when a **parallel** or **serial** construct is encountered.

#### 817 Restrictions

- 818 • See Section 2.17 Fortran Optional Arguments for discussion of Fortran optional arguments in  
819 **firstprivate** clauses.

### 820 2.5.13. **reduction** clause

821 The **reduction** clause is allowed on the **parallel** and **serial** constructs. It specifies a  
822 reduction operator and one or more *vars*. It implies a **copy** data clause for each reduction *var*,  
823 unless a data clause for that variable appears on the compute construct. For each reduction *var*, a  
824 private copy is created for each parallel gang and initialized for that operator. At the end of the  
825 region, the values for each gang are combined using the reduction operator, and the result combined  
826 with the value of the original *var* and stored in the original *var*. If the reduction *var* is an array or  
827 subarray, the array reduction operation is logically equivalent to applying that reduction operation  
828 to each element of the array or subarray individually. If the reduction *var* is a composite variable,

829 the reduction operation is logically equivalent to applying that reduction operation to each member  
830 of the composite variable individually. The reduction result is available after the region.

831 The following table lists the operators that are valid and the initialization values; in each case, the  
832 initialization value will be cast into the data type of the *var*. For **max** and **min** reductions, the  
833 initialization values are the least representable value and the largest representable value for that data  
834 type, respectively. At a minimum, the supported data types include Fortran **logical** as well as  
835 the numerical data types in C (e.g., **\_Bool**, **char**, **int**, **float**, **double**, **float \_Complex**,  
836 **double \_Complex**), C++ (e.g., **bool**, **char**, **wchar\_t**, **int**, **float**, **double**), and Fortran  
837 (e.g., **integer**, **real**, **double precision**, **complex**). However, for each reduction operator,  
838 the supported data types include only the types permitted as operands to the corresponding operator  
839 in the base language where (1) for max and min, the corresponding operator is less-than and (2) for  
840 other operators, the operands and the result are the same type.

C and C++		Fortran	
operator	initialization value	operator	initialization value
<b>+</b>	<b>0</b>	<b>+</b>	<b>0</b>
<b>*</b>	<b>1</b>	<b>*</b>	<b>1</b>
<b>max</b>	least	<b>max</b>	least
<b>min</b>	largest	<b>min</b>	largest
<b>&amp;</b>	<b>~0</b>	<b>iand</b>	all bits on
<b> </b>	<b>0</b>	<b>ior</b>	<b>0</b>
<b>^</b>	<b>0</b>	<b>ieor</b>	<b>0</b>
<b>&amp;&amp;</b>	<b>1</b>	<b>.and.</b>	<b>.true.</b>
<b>  </b>	<b>0</b>	<b>.or.</b>	<b>.false.</b>
		<b>.eqv.</b>	<b>.true.</b>
		<b>.neqv.</b>	<b>.false.</b>

## 842 Restrictions

- 843 • A *var* in a **reduction** clause must be a scalar variable name, a composite variable name,  
844 an array name, an array element, or a subarray (refer to Section 2.7.1).
- 845 • If the reduction *var* is an array element or a subarray, accessing the elements of the array  
846 outside the specified index range results in unspecified behavior.
- 847 • The reduction *var* may not be a member of a composite variable.
- 848 • If the reduction *var* is a composite variable, each member of the composite variable must be  
849 a supported datatype for the reduction operation.
- 850 • See Section 2.17 Fortran Optional Arguments for discussion of Fortran optional arguments in  
851 **reduction** clauses.

### 852 2.5.14. default clause

853 The **default** clause is optional. The **none** argument tells the compiler to require that all variables  
854 used in the compute construct that do not have predetermined data attributes to explicitly appear  
855 in a data clause on the compute construct, a **data** construct that lexically contains the compute

856 construct, or a visible **declare** directive. The **present** argument causes all arrays or composite  
857 variables used in the compute construct that have implicitly determined data attributes to be treated  
858 as if they appeared in a **present** clause.

## 859 2.6. Data Environment

860 This section describes the data attributes for variables. The data attributes for a variable may be  
861 *predetermined*, *implicitly determined*, or *explicitly determined*. Variables with predetermined data  
862 attributes may not appear in a data clause that conflicts with that data attribute. Variables with  
863 implicitly determined data attributes may appear in a data clause that overrides the implicit attribute.  
864 Variables with explicitly determined data attributes are those which appear in a data clause on a  
865 **data** construct, a compute construct, or a **declare** directive.

866 OpenACC supports systems with accelerators that have discrete memory from the host, systems  
867 with accelerators that share memory with the host, as well as systems where an accelerator shares  
868 some memory with the host but also has some discrete memory that is not shared with the host.  
869 In the first case, no data is in shared memory. In the second case, all data is in shared memory.  
870 In the third case, some data may be in shared memory and some data may be in discrete memory,  
871 although a single array or aggregate data structure must be allocated completely in shared or discrete  
872 memory. When a nested OpenACC construct is executed on the device, the default target device for  
873 that construct is the same device on which the encountering accelerator thread is executing. In that  
874 case, the target device shares memory with the encountering thread.

### 875 2.6.1. Variables with Predetermined Data Attributes

876 The loop variable in a C **for** statement or Fortran **do** statement that is associated with a loop  
877 directive is predetermined to be private to each thread that will execute each iteration of the loop.  
878 Loop variables in Fortran **do** statements within a compute construct are predetermined to be private  
879 to the thread that executes the loop.

880 Variables declared in a C block that is executed in *vector-partitioned* mode are private to the thread  
881 associated with each vector lane. Variables declared in a C block that is executed in *worker-*  
882 *partitioned vector-single* mode are private to the worker and shared across the threads associated  
883 with the vector lanes of that worker. Variables declared in a C block that is executed in *worker-*  
884 *single* mode are private to the gang and shared across the threads associated with the workers and  
885 vector lanes of that gang.

886 A procedure called from a compute construct will be annotated as **seq**, **vector**, **worker**, or  
887 **gang**, as described Section 2.15 Procedure Calls in Compute Regions. Variables declared in **seq**  
888 routine are private to the thread that made the call. Variables declared in **vector** routine are private  
889 to the worker that made the call and shared across the threads associated with the vector lanes of  
890 that worker. Variables declared in **worker** or **gang** routine are private to the gang that made the  
891 call and shared across the threads associated with the workers and vector lanes of that gang.

## 892 2.6.2. Data Regions and Data Lifetimes

893 Data in shared memory is accessible from the current device as well as to the local thread. Such  
894 data is available to the accelerator for the lifetime of the variable. Data not in shared memory must  
895 be copied to and from device memory using data constructs, clauses, and API routines. A *data*  
896 *lifetime* is the duration from when the data is first made available to the accelerator until it becomes  
897 unavailable. For data in shared memory, the data lifetime begins when the data is allocated and  
898 ends when it is deallocated; for statically allocated data, the data lifetime begins when the program  
899 begins and does not end. For data not in shared memory, the data lifetime begins when it is made  
900 present and ends when it is no longer present.

901 There are four types of data regions. When the program encounters a **data** construct, it creates a  
902 data region.

903 When the program encounters a compute construct with explicit data clauses or with implicit data  
904 allocation added by the compiler, it creates a data region that has a duration of the compute construct.

905 When the program enters a procedure, it creates an implicit data region that has a duration of the  
906 procedure. That is, the implicit data region is created when the procedure is called, and exited when  
907 the program returns from that procedure invocation. There is also an implicit data region associated  
908 with the execution of the program itself. The implicit program data region has a duration of the  
909 execution of the program.

910 In addition to data regions, a program may create and delete data on the accelerator using **enter**  
911 **data** and **exit data** directives or using runtime API routines. When the program executes  
912 an **enter data** directive, or executes a call to a runtime API **acc\_copyin** or **acc\_create**  
913 routine, each *var* on the directive or the variable on the runtime API argument list will be made live  
914 on accelerator.

## 915 2.6.3. Data Structures with Pointers

916 This section describes the behavior of data structures that contain pointers. A pointer may be a  
917 C or C++ pointer (e.g., **float\***), a Fortran pointer or array pointer (e.g., **real, pointer,**  
918 **dimension(:)**), or a Fortran allocatable (e.g., **real, allocatable, dimension(:)**).

919 When a data object is copied to device memory, the values are copied exactly. If the data is a data  
920 structure that includes a pointer, or is just a pointer, the pointer value copied to device memory  
921 will be the host pointer value. If the pointer target object is also allocated in or copied to device  
922 memory, the pointer itself needs to be updated with the device address of the target object before  
923 dereferencing the pointer in device memory.

924 An *attach* action updates the pointer in device memory to point to the device copy of the data  
925 that the host pointer targets; see Section 2.7.2. For Fortran array pointers and allocatable arrays,  
926 this includes copying any associated descriptor (dope vector) to the device copy of the pointer.  
927 When the device pointer target is deallocated, the pointer in device memory should be restored  
928 to the host value, so it can be safely copied back to host memory. A *detach* action updates the  
929 pointer in device memory to have the same value as the corresponding pointer in local memory;  
930 see Section 2.7.2. The *attach* and *detach* actions are performed by the **copy**, **copyin**, **copyout**,  
931 **create**, **attach**, and **detach** data clauses (Sections 2.7.3-2.7.12), and the **acc\_attach** and  
932 **acc\_detach** runtime API routines (Sections 3.2.34 and 3.2.35). The *attach* and *detach* actions  
933 use attachment counters to determine when the pointer in device memory needs to be updated; see

934 Section 2.6.7.

## 935 2.6.4. Data Construct

936 **Summary** The **data** construct defines *vars* to be allocated in the current device memory for  
 937 the duration of the region, whether data should be copied from local memory to the current device  
 938 memory upon region entry, and copied from device memory to local memory upon region exit.

939 **Syntax** In C and C++, the syntax of the OpenACC **data** construct is

```
#pragma acc data [clause-list] new-line
    structured block
```

940 and in Fortran, the syntax is

```
!$acc data [clause-list]
    structured block
!$acc end data
```

941 where *clause* is one of the following:

```
if( condition )
copy( var-list )
copyin( [readonly:]var-list )
copyout( var-list )
create( var-list )
no_create( var-list )
present( var-list )
deviceptr( var-list )
attach( var-list )
default( none | present )
```

942 **Description** Data will be allocated in the memory of the current device and copied from local  
 943 memory to device memory, or copied back, as required. The data clauses are described in Sec-  
 944 tion 2.7 Data Clauses. Structured reference counters are incremented for data when entering a data  
 945 region, and decremented when leaving the region, as described in Section 2.6.6 Reference Counters.

### 946 **if clause**

947 The **if** clause is optional; when there is no **if** clause, the compiler will generate code to allocate  
 948 space in the current device memory and move data from and to the local memory as required.  
 949 When an **if** clause appears, the program will conditionally allocate memory in and move data to  
 950 and/or from device memory. When the *condition* in the **if** clause evaluates to zero in C or C++, or  
 951 **.false.** in Fortran, no device memory will be allocated, and no data will be moved. When the  
 952 *condition* evaluates to nonzero in C or C++, or **.true.** in Fortran, the data will be allocated and  
 953 moved as specified. At most one **if** clause may appear.



954 **default clause**

955 The **default** clause is optional. If the **default** clause is present, then for each compute construct  
 956 that is lexically contained within the data construct the behavior will be as if a **default** clause with  
 957 the same value appeared on the compute construct, unless a **default** clause already appears on  
 958 the compute construct. At most one **default** clause may appear.

959 **2.6.5. Enter Data and Exit Data Directives**

960 **Summary** An **enter data** directive may be used to define *vars* to be allocated in the current  
 961 device memory for the remaining duration of the program, or until an **exit data** directive that  
 962 deallocates the data. They also tell whether data should be copied from local memory to device  
 963 memory at the **enter data** directive, and copied from device memory to local memory at the  
 964 **exit data** directive. The dynamic range of the program between the **enter data** directive and  
 965 the matching **exit data** directive is the data lifetime for that data.

966 **Syntax** In C and C++, the syntax of the OpenACC **enter data** directive is

```
#pragma acc enter data clause-list new-line
```

967 and in Fortran, the syntax is

```
!$acc enter data clause-list
```

968 where *clause* is one of the following:

```
if( condition )  

async [( int-expr )]  

wait [( int-expr-list )]  

copyin( var-list )  

create( var-list )  

attach( var-list )
```

969 In C and C++, the syntax of the OpenACC **exit data** directive is

```
#pragma acc exit data clause-list new-line
```

970 and in Fortran, the syntax is

```
!$acc exit data clause-list
```

971 where *clause* is one of the following:

```

if( condition )
async [( int-expr )]
wait [( int-expr-list )]
copyout( var-list )
delete( var-list )
detach( var-list )
finalize

```

972 **Description** At an **enter data** directive, data may be allocated in the current device mem-  
 973 ory and copied from local memory to device memory. This action enters a data lifetime for those  
 974 *vars*, and will make the data available for **present** clauses on constructs within the data life-  
 975 time. Dynamic reference counters are incremented for this data, as described in Section 2.6.6  
 976 Reference Counters. Pointers in device memory may be *attached* to point to the corresponding  
 977 device copy of the host pointer target.

978 At an **exit data** directive, data may be copied from device memory to local memory and deal-  
 979 located from device memory. If no **finalize** clause appears, dynamic reference counters are  
 980 decremented for this data. If a **finalize** clause appears, the dynamic reference counters are set  
 981 to zero for this data. Pointers in device memory may be *detached* so as to have the same value as  
 982 the original host pointer.

983 The data clauses are described in Section 2.7 Data Clauses. Reference counting behavior is de-  
 984 scribed in Section 2.6.6 Reference Counters.

#### 985 **if clause**

986 The **if** clause is optional; when there is no **if** clause, the compiler will generate code to allocate or  
 987 deallocate space in the current device memory and move data from and to local memory. When an  
 988 **if** clause appears, the program will conditionally allocate or deallocate device memory and move  
 989 data to and/or from device memory. When the *condition* in the **if** clause evaluates to zero in C or  
 990 C++, or **.false.** in Fortran, no device memory will be allocated or deallocated, and no data will  
 991 be moved. When the *condition* evaluates to nonzero in C or C++, or **.true.** in Fortran, the data  
 992 will be allocated or deallocated and moved as specified.

#### 993 **async clause**

994 The **async** clause is optional; see Section 2.16 Asynchronous Behavior for more information.

#### 995 **wait clause**

996 The **wait** clause is optional; see Section 2.16 Asynchronous Behavior for more information.

#### 997 **finalize clause**

998 The **finalize** clause is allowed on the **exit data** directive and is optional. When no **finalize**  
 999 clause appears, the **exit data** directive will decrement the dynamic reference counters for *vars*

1000 appearing in **copyout** and **delete** clauses, and will decrement the attachment counters for point-  
1001 ers appearing in **detach** clauses. If a **finalize** clause appears, the **exit data** directive will  
1002 set the dynamic reference counters to zero for *vars* appearing in **copyout** and **delete** clauses,  
1003 and will set the attachment counters to zero for pointers appearing in **detach** clauses.

### 1004 2.6.6. Reference Counters

1005 When device memory is allocated for data not in shared memory due to data clauses or OpenACC  
1006 API routine calls, the OpenACC implementation keeps track of that device memory and its relation-  
1007 ship to the corresponding data in host memory.

1008 Each section of device memory will be associated with two *reference counters* per device, a struc-  
1009 tured reference counter and a dynamic reference counter. The structured and dynamic reference  
1010 counters are used to determine when to allocate or deallocate data in device memory. The struc-  
1011 tured reference counter for a block of data keeps track of how many nested data regions have been  
1012 entered for that data. The initial value of the structured reference counter for static data in device  
1013 memory (in a global **declare** directive) is one; for all other data, the initial value is zero. The  
1014 dynamic reference counter for a block of data keeps track of how many dynamic data lifetimes are  
1015 currently active in device memory for that block. The initial value of the dynamic reference counter  
1016 is zero. Data is considered *present* if the sum of the structured and dynamic reference counters is  
1017 greater than zero.

1018 A structured reference counter is incremented when entering each data or compute region that con-  
1019 tain an explicit data clause or implicitly-determined data attributes for that block of memory, and  
1020 is decremented when exiting that region. A dynamic reference counter is incremented for each  
1021 **enter data copyin** or **create** clause, or each **acc\_copyin** or **acc\_create** API routine  
1022 call for that block of memory. The dynamic reference counter is decremented for each **exit data**  
1023 **copyout** or **delete** clause when no **finalize** clause appears, or each **acc\_copyout** or  
1024 **acc\_delete** API routine call for that block of memory. The dynamic reference counter will be  
1025 set to zero with an **exit data copyout** or **delete** clause when a **finalize** clause appears,  
1026 or each **acc\_copyout\_finalize** or **acc\_delete\_finalize** API routine call for the block  
1027 of memory. The reference counters are modified synchronously with the local thread, even if the  
1028 data directives include an **async** clause. When both structured and dynamic reference counters  
1029 reach zero, the data lifetime in device memory for that data ends.

### 1030 2.6.7. Attachment Counter

1031 Since multiple pointers can target the same address, each pointer in device memory is associated  
1032 with an *attachment counter* per device. The *attachment counter* for a pointer is initialized to zero  
1033 when the pointer is allocated in device memory. The *attachment counter* for a pointer is set to one  
1034 whenever the pointer is *attached* to new target address, and incremented whenever an *attach* action  
1035 for that pointer is performed for the same target address. The *attachment counter* is decremented  
1036 whenever a *detach* action occurs for the pointer, and the pointer is *detached* when the *attachment*  
1037 *counter* reaches zero. This is described in more detail in Section 2.7.2 Data Clause Actions.

1038 A pointer in device memory can be assigned a device address in two ways. The pointer can be  
1039 attached to a device address due to data clauses or API routines, as described in Section 2.7.2  
1040 Data Clause Actions, or the pointer can be assigned in a compute region executed on that device.

1041 Unspecified behavior may result if both ways are used for the same pointer.

1042 Pointer members of structs, classes, or derived types in device or host memory can be overwritten  
1043 due to update directives or API routines. It is the user's responsibility to ensure that the pointers  
1044 have the appropriate values before or after the data movement in either direction. The behavior of  
1045 the program is undefined if any of the pointer members are attached when an update of a composite  
1046 variable is performed.

## 1047 2.7. Data Clauses

1048 These data clauses may appear on the **parallel** construct, **kernels** construct, **serial** con-  
1049 struct, **data** construct, the **enter data** and **exit data** directives, and **declare** directives.  
1050 In the descriptions, the *region* is a compute region with a clause appearing on a **parallel**,  
1051 **kernels**, or **serial** construct, a data region with a clause on a **data** construct, or an implicit  
1052 data region with a clause on a **declare** directive. If the **declare** directive appears in a global  
1053 context, the corresponding implicit data region has a duration of the program. The list argument to  
1054 each data clause is a comma-separated collection of *vars*. For all clauses except **deviceptr** and  
1055 **present**, the list argument may include a Fortran *common block* name enclosed within slashes,  
1056 if that *common block* name also appears in a **declare** directive **link** clause. In all cases, the  
1057 compiler will allocate and manage a copy of the *var* in the memory of the current device, creating a  
1058 visible device copy of that *var*, for data not in shared memory.

1059 OpenACC supports accelerators with discrete memories from the local thread. However, if the  
1060 accelerator can access the local memory directly, the implementation may avoid the memory allo-  
1061 cation and data movement and simply share the data in local memory. Therefore, a program that  
1062 uses and assigns data on the host and uses and assigns the same data on the accelerator within a  
1063 data region without update directives to manage the coherence of the two copies may get different  
1064 answers on different accelerators or implementations.

### 1065 Restrictions

- 1066 • Data clauses may not follow a **device\_type** clause.
- 1067 • See Section 2.17 Fortran Optional Arguments for discussion of Fortran optional arguments in  
1068 data clauses.

### 1069 2.7.1. Data Specification in Data Clauses

1070 In C and C++, a subarray is an array name followed by an extended array range specification in  
1071 brackets, with start and length, such as

**AA[2:n]**

1072 If the lower bound is missing, zero is used. If the length is missing and the array has known size, the  
1073 size of the array is used; otherwise the length is required. The subarray **AA[2:n]** means element  
1074 **AA[2]**, **AA[3]**, ..., **AA[2+n-1]**.

1075 In C and C++, a two dimensional array may be declared in at least four ways:

- 1076 • Statically-sized array: `float AA[100][200];`
- 1077 • Pointer to statically sized rows: `typedef float row[200]; row* BB;`
- 1078 • Statically-sized array of pointers: `float* CC[200];`
- 1079 • Pointer to pointers: `float** DD;`

1080 Each dimension may be statically sized, or a pointer to dynamically allocated memory. Each of  
1081 these may be included in a data clause using subarray notation to specify a rectangular array:

- 1082 • `AA[2:n][0:200]`
- 1083 • `BB[2:n][0:m]`
- 1084 • `CC[2:n][0:m]`
- 1085 • `DD[2:n][0:m]`

1086 Multidimensional rectangular subarrays in C and C++ may be specified for any array with any com-  
1087 bination of statically-sized or dynamically-allocated dimensions. For statically sized dimensions,  
1088 all dimensions except the first must specify the whole extent, to preserve the contiguous data re-  
1089 striction, discussed below. For dynamically allocated dimensions, the implementation will allocate  
1090 pointers in device memory corresponding to the pointers in local memory, and will fill in those  
1091 pointers as appropriate.

1092 In Fortran, a subarray is an array name followed by a comma-separated list of range specifications  
1093 in parentheses, with lower and upper bound subscripts, such as

**`arr(1:high, low:100)`**

1094 If either the lower or upper bounds are missing, the declared or allocated bounds of the array, if  
1095 known, are used. All dimensions except the last must specify the whole extent, to preserve the  
1096 contiguous data restriction, discussed below.

## 1097 Restrictions

- 1098 • In Fortran, the upper bound for the last dimension of an assumed-size dummy array must be  
1099 specified.
- 1100 • In C and C++, the length for dynamically allocated dimensions of an array must be explicitly  
1101 specified.
- 1102 • In C and C++, modifying pointers in pointer arrays during the data lifetime, either on the host  
1103 or on the device, may result in undefined behavior.
- 1104 • If a subarray is specified in a data clause, the implementation may choose to allocate memory  
1105 for only that subarray on the accelerator.
- 1106 • In Fortran, array pointers may be specified, but pointer association is not preserved in device  
1107 memory.
- 1108 • Any array or subarray in a data clause, including Fortran array pointers, must be a contiguous  
1109 block of memory, except for dynamic multidimensional C arrays.

- 1110 • In C and C++, if a variable or array of composite type is specified, all the data members of the  
1111 struct or class are allocated and copied, as appropriate. If a composite member is a pointer  
1112 type, the data addressed by that pointer are not implicitly copied.
- 1113 • In Fortran, if a variable or array of composite type is specified, all the members of that derived  
1114 type are allocated and copied, as appropriate. If any member has the **allocatable** or  
1115 **pointer** attribute, the data accessed through that member are not copied.
- 1116 • If an expression is used in a subscript or subarray expression in a clause on a **data** construct,  
1117 the same value is used when copying data at the end of the data region, even if the values of  
1118 variables in the expression change during the data region.

## 1119 2.7.2. Data Clause Actions

1120 Most of the data clauses perform one or more the following actions. The actions test or modify one  
1121 or both of the structured and dynamic reference counters, depending on the directive on which the  
1122 data clause appears.

### 1123 Present Increment Action

1124 A *present increment* action is one of the actions that may be performed for a **present** (Section  
1125 2.7.4), **copy** (Section 2.7.5), **copyin** (Section 2.7.6), **copyout** (Section 2.7.7), **create** (Sec-  
1126 tion 2.7.8), or **no\_create** (Section 2.7.9) clause, or for a call to an **acc\_copyin** (Section 3.2.20)  
1127 or **acc\_create** (Section 3.2.21) API routine. See those sections for details.

1128 A *present increment* action for a *var* occurs only when *var* is already present in device memory.

1129 A *present increment* action for a *var* increments the structured or dynamic reference counter for *var*.

### 1130 Present Decrement Action

1131 A *present decrement* action is one of the actions that may be performed for a **present** (Section  
1132 2.7.4), **copy** (Section 2.7.5), **copyin** (Section 2.7.6), **copyout** (Section 2.7.7), **create** (Sec-  
1133 tion 2.7.8), **no\_create** (Section 2.7.9), or **delete** (Section 2.7.10) clause, or for a call to an  
1134 **acc\_copyout** (Section 3.2.22) or **acc\_delete** (Section 3.2.23) API routine. See those sec-  
1135 tions for details.

1136 A *present decrement* action for a *var* occurs only when *var* is already present in device memory.

1137 A *present decrement* action for a *var* decrements the structured or dynamic reference counter for  
1138 *var*, if its value is greater than zero. If the device memory associated with *var* was mapped to  
1139 the device using **acc\_map\_data**, the dynamic reference count may not be decremented to zero,  
1140 except by a call to **acc\_unmap\_data**. If the reference counter is already zero, its value is left  
1141 unchanged.

### 1142 Create Action

1143 A *create* action is one of the actions that may be performed for a **copyout** (Section 2.7.7) or  
1144 **create** (Section 2.7.8) clause, or for a call to an **acc\_create** API routine (Section 3.2.21). See

1145 those sections for details.

1146 A *create* action for a *var* occurs only when *var* is not already present in device memory.

1147 A *create* action for a *var*:

- 1148 • allocates device memory for *var*; and
- 1149 • sets the structured or dynamic reference counter to one.

### 1150 **Copyin Action**

1151 A *copyin* action is one of the actions that may be performed for a **copy** (Section 2.7.5) or **copyin**  
1152 (Section 2.7.6) clause, or for a call to an **acc\_copyin** API routine (Section 3.2.20). See those  
1153 sections for details.

1154 A *copyin* action for a *var* occurs only when *var* is not already present in device memory.

1155 A *copyin* action for a *var*:

- 1156 • allocates device memory for *var*;
- 1157 • initiates a copy of the data for *var* from the local thread memory to the corresponding device  
1158 memory; and
- 1159 • sets the structured or dynamic reference counter to one.

1160 The data copy may complete asynchronously, depending on other clauses on the directive.

### 1161 **Copyout Action**

1162 A *copyout* action is one of the actions that may be performed for a **copy** (Section 2.7.5) or  
1163 **copyout** (Section 2.7.7) clause, or for a call to an **acc\_copyout** API routine (Section 3.2.22).  
1164 See those sections for details.

1165 A *copyout* action for a *var* occurs only when *var* is present in device memory.

1166 A *copyout* action for a *var*:

- 1167 • performs an *immediate detach* action for any pointer in *var*;
- 1168 • initiates a copy of the data for *var* from device memory to the corresponding local thread  
1169 memory; and
- 1170 • deallocates device memory for *var*.

1171 The data copy may complete asynchronously, depending on other clauses on the directive, in which  
1172 case the memory is deallocated when the data copy is complete.

### 1173 **Delete Action**

1174 A *delete* action is one of the actions that may be performed for a **present** (Section 2.7.4), **copyin**  
1175 (Section 2.7.6), **create** (Section 2.7.8), **no\_create** (Section 2.7.9), or **delete** (Section 2.7.10)  
1176 clause, or for a call to an **acc\_delete** API routine (Section 3.2.23). See those sections for details.

1177 A *delete* action for a *var* occurs only when *var* is present in device memory.

1178 A *delete* action for *var*:

- 1179 • performs an *immediate detach* action for any pointer in *var*; and
- 1180 • deallocates device memory for *var*.

### 1181 **Attach Action**

1182 An *attach* action is one of the actions that may be performed for a **present** (Section 2.7.4),  
1183 **copy** (Section 2.7.5), **copyin** (Section 2.7.6), **copyout** (Section 2.7.7), **create** (Section 2.7.8),  
1184 **no\_create** (Section 2.7.9), or **attach** (Section 2.7.10) clause, or for a call to an **acc\_attach**  
1185 API routine (Section 3.2.34). See those sections for details.

1186 An *attach* action for a *var* occurs only when *var* is a pointer reference.

1187 If the pointer *var* is in shared memory or is not present in the current device memory, or if the  
1188 address to which *var* points is not present in the current device memory, no action is taken. If the  
1189 *attachment counter* for *var* is nonzero and the pointer in device memory already points to the device  
1190 copy of the data in *var*, the *attachment counter* for the pointer *var* is incremented. Otherwise, the  
1191 pointer in device memory is *attached* to the device copy of the data by initiating an update for the  
1192 pointer in device memory to point to the device copy of the data and setting the *attachment counter*  
1193 for the pointer *var* to one. The update may complete asynchronously, depending on other clauses  
1194 on the directive. The pointer update must follow any data copies due to *copyin* actions that are  
1195 performed for the same directive.

### 1196 **Detach Action**

1197 A *detach* action is one of the actions that may be performed for a **present** (Section 2.7.4),  
1198 **copy** (Section 2.7.5), **copyin** (Section 2.7.6), **copyout** (Section 2.7.7), **create** (Section 2.7.8),  
1199 **no\_create** (Section 2.7.9), **delete** (Section 2.7.10), or **detach** (Section 2.7.10) clause, or for  
1200 a call to an **acc\_detach** API routine (Section 3.2.35). See those sections for details.

1201 A *detach* action for a *var* occurs only when *var* is a pointer reference.

1202 If the pointer *var* is in shared memory or is not present in the current device memory, or if the  
1203 *attachment counter* for *var* for the pointer is zero, no action is taken. Otherwise, the *attachment*  
1204 *counter* for the pointer *var* is decremented. If the *attachment counter* is decreased to zero, the  
1205 pointer is *detached* by initiating an update for the pointer *var* in device memory to have the same  
1206 value as the corresponding pointer in local memory. The update may complete asynchronously,  
1207 depending on other clauses on the directive. The pointer update must precede any data copies due  
1208 to *copyout* actions that are performed for the same directive.

### 1209 **Immediate Detach Action**

1210 An *immediate detach* action is one of the actions that may be performed for a **detach** (Section  
1211 2.7.10) clause, or for a call to an **acc\_detach\_finalize** API routine (Section 3.2.35). See  
1212 those sections for details.

1213 An *immediate detach* action for a *var* occurs only when *var* is a pointer reference and is present in  
1214 device memory.



1215 If the *attachment counter* for the pointer is zero, the *immediate detach* action has no effect. Other-  
 1216 wise, the *attachment counter* for the pointer set to zero and the pointer is *detached* by initiating an  
 1217 update for the pointer in device memory to have the same value as the corresponding pointer in local  
 1218 memory. The update may complete asynchronously, depending on other clauses on the directive.  
 1219 The pointer update must precede any data copies due to *copyout* actions that are performed for the  
 1220 same directive.

### 1221 2.7.3. deviceptr clause

1222 The **deviceptr** clause may appear on structured **data** and compute constructs and **declare**  
 1223 directives.

1224 The **deviceptr** clause is used to declare that the pointers in *var-list* are device pointers, so the  
 1225 data need not be allocated or moved between the host and device for this pointer.

1226 In C and C++, the *vars* in *var-list* must be pointer variables.

1227 In Fortran, the *vars* in *var-list* must be dummy arguments (arrays or scalars), and may not have the  
 1228 Fortran **pointer**, **allocatable**, or **value** attributes.

1229 For data in shared memory, host pointers are the same as device pointers, so this clause has no  
 1230 effect.

### 1231 2.7.4. present clause

1232 The **present** clause may appear on structured **data** and compute constructs and **declare** di-  
 1233 rectives. The **present** clause specifies that *vars* in *var-list* are in shared memory or are already  
 1234 present in the current device memory due to data regions or data lifetimes that contain the construct  
 1235 on which the **present** clause appears.

1236 For each *var* in *varlist*, if *var* is in shared memory, no action is taken; if *var* is not in shared memory,  
 1237 the **present** clause behaves as follows:

- 1238 • At entry to the region:
  - 1239 – If *var* is not present in the current device memory, a runtime error is issued.
  - 1240 – Otherwise, a *present increment* action with the structured reference counter is performed.
  - 1241 If *var* is a pointer reference, an *attach* action is performed.
- 1242 • At exit from the region:
  - 1243 – If *var* is not present in the current device memory, a runtime error is issued.
  - 1244 – Otherwise, a *present decrement* action with the structured reference counter is per-  
 1245 formed. If *var* is a pointer reference, a *detach* action is performed. If both structured  
 1246 and dynamic reference counters are zero, a *delete* action is performed.

### 1247 Restrictions

- 1248 • If only a subarray of an array is present in the current device memory, the **present** clause  
 1249 must specify the same subarray, or a subarray that is a proper subset of the subarray in the  
 1250 data lifetime.

- 1251       • It is a runtime error if the subarray in *var-list* clause includes array elements that are not part  
1252       of the subarray specified in the data lifetime.

### 1253 2.7.5. copy clause

1254 The **copy** clause may appear on structured **data** and compute constructs and on **declare** direc-  
1255 tives.

1256 For each *var* in *varlist*, if *var* is in shared memory, no action is taken; if *var* is not in shared memory,  
1257 the **copy** clause behaves as follows:

- 1258       • At entry to the region:
  - 1259           – If *var* is present, a *present increment* action with the structured reference counter is  
1260           performed. If *var* is a pointer reference, an *attach* action is performed.
  - 1261           – Otherwise, a *copyin* action with the structured reference counter is performed. If *var* is  
1262           a pointer reference, an *attach* action is performed.
- 1263       • At exit from the region:
  - 1264           – If *var* is not present in the current device memory, a runtime error is issued.
  - 1265           – Otherwise, a *present decrement* action with the structured reference counter is per-  
1266           formed. If *var* is a pointer reference, a *detach* action is performed. If both structured  
1267           and dynamic reference counters are zero, a *copyout* action is performed.

1268 The restrictions regarding subarrays in the **present** clause apply to this clause.

1269 For compatibility with OpenACC 2.0, **present\_or\_copy** and **pcopy** are alternate names for  
1270 **copy**.

### 1271 2.7.6. copyin clause

1272 The **copyin** clause may appear on structured **data** and compute constructs, on **declare** direc-  
1273 tives, and on **enter data** directives.

1274 For each *var* in *varlist*, if *var* is in shared memory, no action is taken; if *var* is not in shared memory,  
1275 the **copyin** clause behaves as follows:

- 1276       • At entry to a region, the structured reference counter is used. On an **enter data** directive,  
1277       the dynamic reference counter is used.
  - 1278           – If *var* is present, a *present increment* action with the appropriate reference counter is  
1279           performed. If *var* is a pointer reference, an *attach* action is performed.
  - 1280           – Otherwise, a *copyin* action with the appropriate reference counter is performed. If *var*  
1281           is a pointer reference, an *attach* action is performed.
- 1282       • At exit from the region:
  - 1283           – If *var* is not present in the current device memory, a runtime error is issued.
  - 1284           – Otherwise, a *present decrement* action with the structured reference counter is per-  
1285           formed. If *var* is a pointer reference, a *detach* action is performed. If both structured  
1286           and dynamic reference counters are zero, a *delete* action is performed.

1287 If the optional **readonly** modifier appears, then the implementation may assume that the data  
 1288 referenced by *var-list* is never written to within the applicable region.

1289 The restrictions regarding subarrays in the **present** clause apply to this clause.

1290 For compatibility with OpenACC 2.0, **present\_or\_copyin** and **pcopyin** are alternate names  
 1291 for **copyin**.

1292 An **enter data** directive with a **copyin** clause is functionally equivalent to a call to the **acc\_copyin**  
 1293 API routine, as described in Section 3.2.20.

### 1294 2.7.7. copyout clause

1295 The **copyout** clause may appear on structured **data** and compute constructs, on **declare** di-  
 1296 rectives, and on **exit data** directives.

1297 For each *var* in *varlist*, if *var* is in shared memory, no action is taken; if *var* is not in shared memory,  
 1298 the **copyout** clause behaves as follows:

- 1299 • At entry to a region:
    - 1300 – If *var* is present, a *present increment* action with the structured reference counter is  
 1301 performed. If *var* is a pointer reference, an *attach* action is performed.
    - 1302 – Otherwise, a *create* action with the structured reference is performed. If *var* is a pointer  
 1303 reference, an *attach* action is performed.
  - 1304 • At exit from a region, the structured reference counter is used. On an **exit data** directive,  
 1305 the dynamic reference counter is used.
    - 1306 – If *var* is not present in the current device memory, a runtime error is issued.
    - 1307 – Otherwise, the reference counter is updated:
      - 1308 \* On an **exit data** directive with a **finalize** clause, the dynamic reference  
 1309 counter is set to zero.
      - 1310 \* Otherwise, a *present decrement* action with the appropriate reference counter is  
 1311 performed.
- 1312 If *var* is a pointer reference, a *detach* action is performed. If both structured and dynamic  
 1313 reference counters are zero, a *copyout* action is performed.

1314 The restrictions regarding subarrays in the **present** clause apply to this clause.

1315 For compatibility with OpenACC 2.0, **present\_or\_copyout** and **pcopyout** are alternate  
 1316 names for **copyout**.

1317 An **exit data** directive with a **copyout** clause and with or without a **finalize** clause is func-  
 1318 tionally equivalent to a call to the **acc\_copyout\_finalize** or **acc\_copyout** API routine,  
 1319 respectively, as described in Section 3.2.22.

### 1320 2.7.8. create clause

1321 The **create** clause may appear on structured **data** and compute constructs, on **declare** direc-  
 1322 tives, and on **enter data** directives.

1323 For each *var* in *varlist*, if *var* is in shared memory, no action is taken; if *var* is not in shared memory,  
1324 the **create** clause behaves as follows:

- 1325 • At entry to a region, the structured reference counter is used. On an **enter data** directive,  
1326 the dynamic reference counter is used.
  - 1327 – If *var* is present, a *present increment* action with the appropriate reference counter is  
1328 performed. If *var* is a pointer reference, an *attach* action is performed.
  - 1329 – Otherwise, a *create* action with the appropriate reference counter is performed. If *var* is  
1330 a pointer reference, an *attach* action is performed.
- 1331 • At exit from the region:
  - 1332 – If *var* is not present in the current device memory, a runtime error is issued.
  - 1333 – Otherwise, a *present decrement* action with the structured reference counter is per-  
1334 formed. If *var* is a pointer reference, a *detach* action is performed. If both structured  
1335 and dynamic reference counters are zero, a *delete* action is performed.

1336 The restrictions regarding subarrays in the **present** clause apply to this clause.

1337 For compatibility with OpenACC 2.0, **present\_or\_create** and **pcreate** are alternate names  
1338 for **create**.

1339 An **enter data** directive with a **create** clause is functionally equivalent to a call to the **acc\_create**  
1340 API routine, as described in Section 3.2.21.

### 1341 2.7.9. no\_create clause

1342 The **no\_create** clause may appear on structured **data** and compute constructs.

1343 For each *var* in *varlist*, if *var* is in shared memory, no action is taken; if *var* is not in shared memory,  
1344 the **no\_create** clause behaves as follows:

- 1345 • At entry to the region:
  - 1346 – If *var* is present, a *present increment* action with the structured reference counter is  
1347 performed. If *var* is a pointer reference, an *attach* action is performed.
  - 1348 – Otherwise, no action is performed, and any device code in this construct will use the  
1349 local memory address for *var*.
- 1350 • At exit from the region:
  - 1351 – If *var* is not present in the current device memory, no action is performed.
  - 1352 – Otherwise, a *present decrement* action with the structured reference counter is per-  
1353 formed. If *var* is a pointer reference, a *detach* action is performed. If both structured  
1354 and dynamic reference counters are zero, a *delete* action is performed.

1355 The restrictions regarding subarrays in the **present** clause apply to this clause.

### 1356 2.7.10. delete clause

1357 The **delete** clause may appear on **exit data** directives.

1358 For each *var* in *varlist*, if *var* is in shared memory, no action is taken; if *var* is not in shared memory,  
1359 the **delete** clause behaves as follows:

- 1360 • If *var* is not present in the current device memory, a runtime error is issued.
  - 1361 • Otherwise, the dynamic reference counter is updated:
    - 1362 – On an **exit data** directive with a **finalize** clause, the dynamic reference counter  
1363 is set to zero.
    - 1364 – Otherwise, a *present decrement* action with the dynamic reference counter is performed.
- 1365 If *var* is a pointer reference, a *detach* action is performed. If both structured and dynamic  
1366 reference counters are zero, a *delete* action is performed.

1367 An **exit data** directive with a **delete** clause and with or without a **finalize** clause is func-  
1368 tionally equivalent to a call to the **acc\_delete\_finalize** or **acc\_delete** API routine, re-  
1369 spectively, as described in Section 3.2.23.

### 1370 2.7.11. attach clause

1371 The **attach** clause may appear on structured **data** and compute constructs and on **enter data**  
1372 directives. Each *var* argument to an **attach** clause must be a C or C++ pointer or a Fortran variable  
1373 or array with the **pointer** or **allocatable** attribute.

1374 For each *var* in *varlist*, if *var* is in shared memory, no action is taken; if *var* is not in shared memory,  
1375 the **attach** clause behaves as follows:

- 1376 • At entry to a region or at an **enter data** directive, an *attach* action is performed.
- 1377 • At exit from the region, a *detach* action is performed.

### 1378 2.7.12. detach clause

1379 The **detach** clause may appear on **exit data** directives. Each *var* argument to a **detach** clause  
1380 must be a C or C++ pointer or a Fortran variable or array with the **pointer** or **allocatable**  
1381 attribute.

1382 For each *var* in *varlist*, if *var* is in shared memory, no action is taken; if *var* is not in shared memory,  
1383 the **detach** clause behaves as follows:

- 1384 • If there is a **finalize** clause on the **exit data** directive, an *immediate detach* action is  
1385 performed.
- 1386 • Otherwise, a *detach* action is performed.

## 1387 2.8. Host\_Data Construct

1388 **Summary** The **host\_data** construct makes the address of data in device memory available on  
1389 the host.

1390 **Syntax** In C and C++, the syntax of the OpenACC **host\_data** construct is

```
#pragma acc host_data clause-list new-line
    structured block
```

1391 and in Fortran, the syntax is

```
!$acc host_data clause-list
    structured block
!$acc end host_data
```

1392 where *clause* is one of the following:

```
use_device( var-list )
if( condition )
if_present
```

1393 **Description** This construct is used to make the address of data in device memory available in  
1394 host code.

### 1395 Restrictions

- 1396 • A *var* in a **use\_device** clause must be the name of a variable or array.
- 1397 • At least one **use\_device** clause must appear.
- 1398 • At most one **if** clause may appear. In Fortran, the condition must evaluate to a scalar logical  
1399 value; in C or C++, the condition must evaluate to a scalar integer value.
- 1400 • See Section 2.17 Fortran Optional Arguments for discussion of Fortran optional arguments in  
1401 **use\_device** clauses.

#### 1402 2.8.1. use\_device clause

1403 The **use\_device** clause tells the compiler to use the current device address of any *var* in *var-list*  
1404 in code within the construct. In particular, this may be used to pass the device address of *var* to  
1405 optimized procedures written in a lower-level API. When there is no **if\_present** clause, and  
1406 either there is no **if** clause or the condition in the **if** clause evaluates to nonzero (in C or C++)  
1407 or **.true.** (in Fortran), the *var* in *var-list* must be present in the accelerator memory due to data  
1408 regions or data lifetimes that contain this construct. For data in shared memory, the device address  
1409 is the same as the host address.

#### 1410 2.8.2. if clause

1411 The **if** clause is optional. When an **if** clause appears and the condition evaluates to zero in C  
1412 or C++, or **.false.** in Fortran, the compiler will not replace the addresses of any *var* in code  
1413 within the construct. When there is no **if** clause, or when an **if** clause appears and the condition  
1414 evaluates to nonzero in C or C++, or **.true.** in Fortran, the compiler will replace the addresses as  
1415 described in the previous subsection.

1416 **2.8.3. if\_present clause**

1417 When an **if\_present** clause appears on the directive, the compiler will only replace the address  
 1418 of any *var* which appears in *var-list* that is present in the current device memory.

1419 **2.9. Loop Construct**

1420 **Summary** The OpenACC **loop** construct applies to a loop which must immediately follow this  
 1421 directive. The **loop** construct can describe what type of parallelism to use to execute the loop and  
 1422 declare private *vars* and reduction operations.

1423 **Syntax** In C and C++, the syntax of the **loop** construct is

```
#pragma acc loop [clause-list] new-line
           for loop
```

1424 In Fortran, the syntax of the **loop** construct is

```
!$acc loop [clause-list]
           do loop
```

1425 where *clause* is one of the following:

```
collapse( n )
gang [ ( gang-arg-list ) ]
worker [ ( [num:]int-expr ) ]
vector [ ( [length:]int-expr ) ]
seq
auto
tile( size-expr-list )
device_type( device-type-list )
independent
private( var-list )
reduction( operator:var-list )
```

1426 where *gang-arg* is one of:

```
[num:]int-expr
static:size-expr
```

1427 and *gang-arg-list* may have at most one **num** and one **static** argument,

1428 and where *size-expr* is one of:

\*  
*int-expr*

1429 Some clauses are only valid in the context of a **kernels** construct; see the descriptions below.

1430 An *orphaned* **loop** construct is a **loop** construct that is not lexically enclosed within a compute  
1431 construct. The parent compute construct of a **loop** construct is the nearest compute construct that  
1432 lexically contains the **loop** construct.

### 1433 Restrictions

- 1434 • Only the **collapse**, **gang**, **worker**, **vector**, **seq**, **auto**, and **tile** clauses may follow  
1435 a **device\_type** clause.
- 1436 • The *int-expr* argument to the **worker** and **vector** clauses must be invariant in the kernels  
1437 region.
- 1438 • A loop associated with a **loop** construct that does not have a **seq** clause must be written  
1439 such that the loop iteration count is computable when entering the **loop** construct.

#### 1440 2.9.1. collapse clause

1441 The **collapse** clause is used to specify how many tightly nested loops are associated with the  
1442 **loop** construct. The argument to the **collapse** clause must be a constant positive integer expres-  
1443 sion. If no **collapse** clause appears, only the immediately following loop is associated with the  
1444 **loop** construct.

1445 If more than one loop is associated with the **loop** construct, the iterations of all the associated loops  
1446 are all scheduled according to the rest of the clauses. The trip count for all loops associated with the  
1447 **collapse** clause must be computable and invariant in all the loops.

1448 It is implementation-defined whether a **gang**, **worker** or **vector** clause on the construct is ap-  
1449 plied to each loop, or to the linearized iteration space.

#### 1450 2.9.2. gang clause

1451 When the parent compute construct is a **parallel** construct, or on an orphaned **loop** construct,  
1452 the **gang** clause specifies that the iterations of the associated loop or loops are to be executed in  
1453 parallel by distributing the iterations among the gangs created by the **parallel** construct. A  
1454 **loop** construct with the **gang** clause transitions a compute region from gang-redundant mode  
1455 to gang-partitioned mode. The number of gangs is controlled by the **parallel** construct; only  
1456 the **static** argument is allowed. The loop iterations must be data independent, except for *vars*  
1457 specified in a **reduction** clause. The region of a loop with the **gang** clause may not contain  
1458 another loop with the **gang** clause unless within a nested compute region.

1459 When the parent compute construct is a **kernels** construct, the **gang** clause specifies that the  
1460 iterations of the associated loop or loops are to be executed in parallel across the gangs. An argument  
1461 with no keyword or with the **num** keyword is allowed only when the **num\_gangs** does not appear  
1462 on the **kernels** construct. If an argument with no keyword or an argument after the **num** keyword



1463 is specified, it specifies how many gangs to use to execute the iterations of this loop. The region of a  
1464 loop with the **gang** clause may not contain another loop with a **gang** clause unless within a nested  
1465 compute region.

1466 The scheduling of loop iterations to gangs is not specified unless the **static** argument appears as  
1467 an argument. If the **static** argument appears with an integer expression, that expression is used  
1468 as a *chunk* size. If the **static** argument appears with an asterisk, the implementation will select a  
1469 *chunk* size. The iterations are divided into chunks of the selected *chunk* size, and the chunks are  
1470 assigned to gangs starting with gang zero and continuing in round-robin fashion. Two **gang** loops  
1471 in the same parallel region with the same number of iterations, and with **static** clauses with the  
1472 same argument, will assign the iterations to gangs in the same manner. Two **gang** loops in the  
1473 same kernels region with the same number of iterations, the same number of gangs to use, and with  
1474 **static** clauses with the same argument, will assign the iterations to gangs in the same manner.

### 1475 2.9.3. worker clause

1476 When the parent compute construct is a **parallel** construct, or on an orphaned **loop** construct,  
1477 the **worker** clause specifies that the iterations of the associated loop or loops are to be executed  
1478 in parallel by distributing the iterations among the multiple workers within a single gang. A **loop**  
1479 construct with a **worker** clause causes a gang to transition from worker-single mode to worker-  
1480 partitioned mode. In contrast to the **gang** clause, the **worker** clause first activates additional  
1481 worker-level parallelism and then distributes the loop iterations across those workers. No argument  
1482 is allowed. The loop iterations must be data independent, except for *vars* specified in a **reduction**  
1483 clause. The region of a loop with the **worker** clause may not contain a loop with the **gang** or  
1484 **worker** clause unless within a nested compute region.

1485 When the parent compute construct is a **kernels** construct, the **worker** clause specifies that the  
1486 iterations of the associated loop or loops are to be executed in parallel across the workers within  
1487 a single gang. An argument is allowed only when the **num\_workers** does not appear on the  
1488 **kernels** construct. The optional argument specifies how many workers per gang to use to execute  
1489 the iterations of this loop. The region of a loop with the **worker** clause may not contain a loop  
1490 with a **gang** or **worker** clause unless within a nested compute region.

1491 All workers will complete execution of their assigned iterations before any worker proceeds beyond  
1492 the end of the loop.

### 1493 2.9.4. vector clause

1494 When the parent compute construct is a **parallel** construct, or on an orphaned **loop** construct,  
1495 the **vector** clause specifies that the iterations of the associated loop or loops are to be executed in  
1496 vector or SIMD mode. A **loop** construct with a **vector** clause causes a worker to transition from  
1497 vector-single mode to vector-partitioned mode. Similar to the **worker** clause, the **vector** clause  
1498 first activates additional vector-level parallelism and then distributes the loop iterations across those  
1499 vector lanes. The operations will execute using vectors of the length specified or chosen for the  
1500 parallel region. The region of a loop with the **vector** clause may not contain a loop with the  
1501 **gang**, **worker**, or **vector** clause unless within a nested compute region.

1502 When the parent compute construct is a **kernels** construct, the **vector** clause specifies that the  
1503 iterations of the associated loop or loops are to be executed with vector or SIMD processing. An

1504 argument is allowed only when the **vector\_length** does not appear on the **kernels** construct.  
1505 If an argument is specified, the iterations will be processed in vector strips of that length; if no  
1506 argument is specified, the implementation will choose an appropriate vector length. The region of  
1507 a loop with the **vector** clause may not contain a loop with a **gang**, **worker**, or **vector** clause  
1508 unless within a nested compute region.

1509 All vector lanes will complete execution of their assigned iterations before any vector lane proceeds  
1510 beyond the end of the loop.

### 1511 2.9.5. seq clause

1512 The **seq** clause specifies that the associated loop or loops are to be executed sequentially by the  
1513 accelerator. This clause will override any automatic parallelization or vectorization.

### 1514 2.9.6. auto clause

1515 The **auto** clause specifies that the implementation must analyze the loop and determine whether  
1516 the loop iterations are data independent and, if so, select whether to apply parallelism to this loop  
1517 or whether to run the loop sequentially. The implementation may be restricted to the types of  
1518 parallelism it can apply by the presence of **loop** constructs with **gang**, **worker**, or **vector**  
1519 clauses for outer or inner loops. When the parent compute construct is a **kernels** construct, a  
1520 **loop** construct with no **independent** or **seq** clause is treated as if it has the **auto** clause.

### 1521 2.9.7. tile clause

1522 The **tile** clause specifies that the implementation should split each loop in the loop nest into two  
1523 loops, with an outer set of *tile* loops and an inner set of *element* loops. The argument to the **tile**  
1524 clause is a list of one or more tile sizes, where each tile size is a constant positive integer expression  
1525 or an asterisk. If there are  $n$  tile sizes in the list, the **loop** construct must be immediately followed  
1526 by  $n$  tightly-nested loops. The first argument in the *size-expr-list* corresponds to the innermost loop  
1527 of the  $n$  associated loops, and the last element corresponds to the outermost associated loop. If the  
1528 tile size is specified with an asterisk, the implementation will choose an appropriate value. Each  
1529 loop in the nest will be split or *strip-mined* into two loops, an outer *tile* loop and an inner *element*  
1530 loop. The trip count of the element loop will be limited to the corresponding tile size from the  
1531 *size-expr-list*. The *tile* loops will be reordered to be outside all the *element* loops, and the *element*  
1532 loops will all be inside the *tile* loops.

1533 If the **vector** clause appears on the **loop** construct, the **vector** clause is applied to the *element*  
1534 loops. If the **gang** clause appears on the **loop** construct, the **gang** clause is applied to the *tile*  
1535 loops. If the **worker** clause appears on the **loop** construct, the **worker** clause is applied to the  
1536 *element* loops if no **vector** clause appears, and to the *tile* loops otherwise.

### 1537 2.9.8. device\_type clause

1538 The **device\_type** clause is described in Section 2.4 Device-Specific Clauses.

### 1539 2.9.9. independent clause

1540 The **independent** clause tells the implementation that the iterations of this loop are data-independent  
1541 with respect to each other. This allows the implementation to generate code to execute the iterations  
1542 in parallel with no synchronization. When the parent compute construct is a **parallel** construct,  
1543 the **independent** clause is implied on all **loop** constructs without a **seq** or **auto** clause.

#### 1544 Note

- 1545 • It is likely a programming error to use the **independent** clause on a loop if any iteration  
1546 writes to a variable or array element that any other iteration also writes or reads, except for  
1547 *vars* in a **reduction** clause or accesses in atomic regions.

### 1548 2.9.10. private clause

1549 The **private** clause on a **loop** construct specifies that a copy of each item in *var-list* will be  
1550 created. If the body of the loop is executed in *vector-partitioned* mode, a copy of the item is created  
1551 for each thread associated with each vector lane. If the body of the loop is executed in *worker-*  
1552 *partitioned vector-single* mode, a copy of the item is created for and shared across the set of threads  
1553 associated with all the vector lanes of each worker. Otherwise, a copy of the item is created for and  
1554 shared across the set of threads associated with all the vector lanes of all the workers of each gang.

#### 1555 Restrictions

- 1556 • See Section 2.17 Fortran Optional Arguments for discussion of Fortran optional arguments in  
1557 **private** clauses.

### 1558 2.9.11. reduction clause

1559 The **reduction** clause specifies a reduction operator and one or more *vars*. For each reduction  
1560 *var*, a private copy is created in the same manner as for a **private** clause on the **loop** construct,  
1561 and initialized for that operator; see the table in Section 2.5.13 reduction clause. At the end of the  
1562 loop, the values for each thread are combined using the specified reduction operator, and the result  
1563 combined with the value of the original *var* and stored in the original *var* at the end of the parallel  
1564 or kernels region if the loop has gang parallelism, and at the end of the loop otherwise. If the  
1565 reduction *var* is an array or subarray, the reduction operation is logically equivalent to applying that  
1566 reduction operation to each array element of the array or subarray individually. If the reduction *var*  
1567 is a composite variable, the reduction operation is logically equivalent to applying that reduction  
1568 operation to each member of the composite variable individually.

1569 In a parallel region, if the **reduction** clause is used on a loop with the **vector** or **worker**  
1570 clauses (and no **gang** clause), and the reduction *var* is private, the value of the private reduction  
1571 *var* will be updated at the exit of the loop. If the reduction *var* is not private in the parallel region,  
1572 or if the **reduction** clause is used on a loop with the **gang** clause, the value of the reduction *var*  
1573 will not be updated until the end of the parallel region.

1574 If a variable is involved in a reduction that spans multiple nested loops where two or more of those  
1575 loops have associated **loop** directives, a **reduction** clause containing that variable must appear

1576 on each of those **loop** directives.

## 1577 Restrictions

- 1578 • A *var* in a **reduction** clause must be a scalar variable name, a composite variable name,  
1579 an array name, an array element, or a subarray (refer to Section 2.7.1).
- 1580 • Reduction clauses on nested constructs for the same reduction *var* must have the same reduc-  
1581 tion operator.
- 1582 • The **reduction** clause may not appear on an orphaned **loop** construct with the **gang**  
1583 clause, or on an orphaned **loop** construct that will generate gang parallelism in a procedure  
1584 that is compiled with the **routine gang** clause.
- 1585 • The restrictions for a **reduction** clause on a compute construct listed in in Section 2.5.13  
1586 reduction clause also apply to a **reduction** clause on a loop construct.
- 1587 • See Section 2.17 Fortran Optional Arguments for discussion of Fortran optional arguments in  
1588 **reduction** clauses.

## 1589 2.10. Cache Directive

1590 **Summary** The **cache** directive may appear at the top of (inside of) a loop. It specifies array  
1591 elements or subarrays that should be fetched into the highest level of the cache for the body of the  
1592 loop.

1593 **Syntax** In C and C++, the syntax of the **cache** directive is

```
#pragma acc cache( [readonly:]var-list ) new-line
```

1594 In Fortran, the syntax of the **cache** directive is

```
!$acc cache( [readonly:]var-list )
```

1595 A *var* in a **cache** directive must be a single array element or a simple subarray. In C and C++,  
1596 a simple subarray is an array name followed by an extended array range specification in brackets,  
1597 with start and length, such as

```
arr [lower:length]
```

1598 where the lower bound is a constant, loop invariant, or the **for** loop index variable plus or minus a  
1599 constant or loop invariant, and the length is a constant.

1600 In Fortran, a simple subarray is an array name followed by a comma-separated list of range specifi-  
1601 cations in parentheses, with lower and upper bound subscripts, such as

```
arr (lower:upper, lower2:upper2)
```

1602 The lower bounds must be constant, loop invariant, or the **do** loop index variable plus or minus  
 1603 a constant or loop invariant; moreover the difference between the corresponding upper and lower  
 1604 bounds must be a constant.

1605 If the optional **readonly** modifier appears, then the implementation may assume that the data  
 1606 referenced by any *var* in that directive is never written to within the applicable region.

### 1607 Restrictions

- 1608 • If an array element or subarray is listed in a **cache** directive, all references to that array  
 1609 during execution of that loop iteration must not refer to elements of the array outside the  
 1610 index range specified in the **cache** directive.
- 1611 • See Section 2.17 Fortran Optional Arguments for discussion of Fortran optional arguments in  
 1612 **cache** directives.

## 1613 2.11. Combined Constructs

1614 **Summary** The combined OpenACC **parallel loop**, **kernels loop**, and **serial loop**  
 1615 constructs are shortcuts for specifying a **loop** construct nested immediately inside a **parallel**,  
 1616 **kernels**, or **serial** construct. The meaning is identical to explicitly specifying a **parallel**,  
 1617 **kernels**, or **serial** construct containing a **loop** construct. Any clause that is allowed on a  
 1618 **parallel** or **loop** construct is allowed on the **parallel loop** construct; any clause allowed  
 1619 on a **kernels** or **loop** construct is allowed on a **kernels loop** construct; and any clause  
 1620 allowed on a **serial** or **loop** construct is allowed on a **serial loop** construct.

1621 **Syntax** In C and C++, the syntax of the **parallel loop** construct is

```
#pragma acc parallel loop [clause-list] new-line
    for loop
```

1622 In Fortran, the syntax of the **parallel loop** construct is

```
!$acc parallel loop [clause-list]
    do loop
[!$acc end parallel loop]
```

1623 The associated structured block is the loop which must immediately follow the directive. Any of  
 1624 the **parallel** or **loop** clauses valid in a parallel region may appear.

1625 In C and C++, the syntax of the **kernels loop** construct is

```
#pragma acc kernels loop [clause-list] new-line
    for loop
```

1626 In Fortran, the syntax of the **kernels loop** construct is

```

!$acc kernels loop [clause-list]
    do loop
!$acc end kernels loop

```

1627 The associated structured block is the loop which must immediately follow the directive. Any of  
 1628 the **kernels** or **loop** clauses valid in a kernels region may appear.

1629 In C and C++, the syntax of the **serial loop** construct is

```

#pragma acc serial loop [clause-list] new-line
    for loop

```

1630 In Fortran, the syntax of the **serial loop** construct is

```

!$acc serial loop [clause-list]
    do loop
!$acc end serial loop

```

1631 The associated structured block is the loop which must immediately follow the directive. Any of  
 1632 the **serial** or **loop** clauses valid in a serial region may appear.

1633 A **private** or **reduction** clause on a combined construct is treated as if it appeared on the  
 1634 **loop** construct. In addition, a **reduction** clause on a combined construct implies a **copy** data  
 1635 clause for each reduction variable, unless a data clause for that variable appears on the combined  
 1636 construct.

### 1637 Restrictions

- 1638 • The restrictions for the **parallel**, **kernels**, **serial**, and **loop** constructs apply.

## 1639 2.12. Atomic Construct

1640 **Summary** An **atomic** construct ensures that a specific storage location is accessed and/or up-  
 1641 dated atomically, preventing simultaneous reading and writing by gangs, workers, and vector threads  
 1642 that could result in indeterminate values.

1643 **Syntax** In C and C++, the syntax of the **atomic** constructs is:

```

#pragma acc atomic [atomic-clause] new-line
    expression-stmt

```

1644 OR:

```

#pragma acc atomic update capture new-line
    structured-block

```



1652 In the preceding expressions:

- 1653 • **x** and **v** (as applicable) are both l-value expressions with scalar type.
- 1654 • During the execution of an atomic region, multiple syntactic occurrences of **x** must designate  
1655 the same storage location.
- 1656 • Neither of **v** and *expr* (as applicable) may access the storage location designated by **x**.
- 1657 • Neither of **x** and *expr* (as applicable) may access the storage location designated by **v**.
- 1658 • *expr* is an expression with scalar type.
- 1659 • *binop* is one of **+**, **\***, **-**, **/**, **&**, **^**, **|**, **<<**, or **>>**.
- 1660 • *binop*, *binop=*, **++**, and **--** are not overloaded operators.
- 1661 • The expression **x binop expr** must be mathematically equivalent to **x binop (expr)**. This  
1662 requirement is satisfied if the operators in *expr* have precedence greater than *binop*, or by  
1663 using parentheses around *expr* or subexpressions of *expr*.
- 1664 • The expression *expr binop x* must be mathematically equivalent to *(expr) binop x*. This  
1665 requirement is satisfied if the operators in *expr* have precedence equal to or greater than *binop*,  
1666 or by using parentheses around *expr* or subexpressions of *expr*.
- 1667 • For forms that allow multiple occurrences of **x**, the number of times that **x** is evaluated is  
1668 unspecified.

1669 In Fortran the syntax of the **atomic** constructs is:

```
!$acc atomic read
  capture-statement
!$acc end atomic
```

1670 OR

```
!$acc atomic write
  write-statement
!$acc end atomic
```

1671 OR

```
!$acc atomic [update]
  update-statement
!$acc end atomic
```

1672 OR

```
!$acc atomic capture
  update-statement
  capture-statement
!$acc end atomic
```



1673 OR

```

!$acc atomic capture
  capture-statement
  update-statement
!$acc end atomic

```

1674 OR

```

!$acc atomic capture
  capture-statement
  write-statement
!$acc end atomic

```

1675 where *write-statement* has the following form (if *atomic-clause* is **write** or **capture**):

```
x = expr
```

1676 where *capture-statement* has the following form (if *atomic-clause* is **capture** or **read**):

```
v = x
```

1677 and where *update-statement* has one of the following forms (if *atomic-clause* is **update**, **capture**,  
1678 or no clause appears):

```

x = x operator expr
x = expr operator x
x = intrinsic_procedure_name ( x, expr-list )
x = intrinsic_procedure_name ( expr-list, x )

```

1679 In the preceding statements:

- 1680 ● **x** and **v** (as applicable) are both scalar variables of intrinsic type.
- 1681 ● **x** must not be an allocatable variable.
- 1682 ● During the execution of an atomic region, multiple syntactic occurrences of **x** must designate  
1683 the same storage location.
- 1684 ● None of **v**, *expr*, and *expr-list* (as applicable) may access the same storage location as **x**.
- 1685 ● None of **x**, *expr*, and *expr-list* (as applicable) may access the same storage location as **v**.
- 1686 ● *expr* is a scalar expression.
- 1687 ● *expr-list* is a comma-separated, non-empty list of scalar expressions. If *intrinsic\_procedure\_name*  
1688 refers to **iand**, **ior**, or **ieor**, exactly one expression must appear in *expr-list*.

- 1689 • *intrinsic\_procedure\_name* is one of **max**, **min**, **iand**, **ior**, or **ieor**. *operator* is one of **+**,
- 1690 **\***, **-**, **/**, **.and.**, **.or.**, **.eqv.**, or **.neqv.**
- 1691 • The expression **x operator expr** must be mathematically equivalent to **x operator (expr)**.
- 1692 This requirement is satisfied if the operators in *expr* have precedence greater than *operator*,
- 1693 or by using parentheses around *expr* or subexpressions of *expr*.
- 1694 • The expression **expr operator x** must be mathematically equivalent to **(expr) operator x**.
- 1695 This requirement is satisfied if the operators in *expr* have precedence equal to or greater than
- 1696 *operator*, or by using parentheses around *expr* or subexpressions of *expr*.
- 1697 • *intrinsic\_procedure\_name* must refer to the intrinsic procedure name and not to other program
- 1698 entities.
- 1699 • *operator* must refer to the intrinsic operator and not to a user-defined operator. All assign-
- 1700 ments must be intrinsic assignments.
- 1701 • For forms that allow multiple occurrences of **x**, the number of times that **x** is evaluated is
- 1702 unspecified.

1703 An **atomic** construct with the **read** clause forces an atomic read of the location designated by **x**.  
 1704 An **atomic** construct with the **write** clause forces an atomic write of the location designated by  
 1705 **x**.

1706 An **atomic** construct with the **update** clause forces an atomic update of the location designated  
 1707 by **x** using the designated operator or intrinsic. Note that when no clause appears, the semantics  
 1708 are equivalent to **atomic update**. Only the read and write of the location designated by **x** are  
 1709 performed mutually atomically. The evaluation of *expr* or *expr-list* need not be atomic with respect  
 1710 to the read or write of the location designated by **x**.

1711 An **atomic** construct with the **capture** clause forces an atomic update of the location designated  
 1712 by **x** using the designated operator or intrinsic while also capturing the original or final value of  
 1713 the location designated by **x** with respect to the atomic update. The original or final value of the  
 1714 location designated by **x** is written into the location designated by **v** depending on the form of the  
 1715 **atomic** construct structured block or statements following the usual language semantics. Only  
 1716 the read and write of the location designated by **x** are performed mutually atomically. Neither the  
 1717 evaluation of *expr* or *expr-list*, nor the write to the location designated by **v**, need to be atomic with  
 1718 respect to the read or write of the location designated by **x**.

1719 For all forms of the **atomic** construct, any combination of two or more of these **atomic** constructs  
 1720 enforces mutually exclusive access to the locations designated by **x**. To avoid race conditions, all  
 1721 accesses of the locations designated by **x** that could potentially occur in parallel must be protected  
 1722 with an **atomic** construct.

1723 Atomic regions do not guarantee exclusive access with respect to any accesses outside of atomic re-  
 1724 gions to the same storage location **x** even if those accesses occur during the execution of a reduction  
 1725 clause.

1726 If the storage location designated by **x** is not size-aligned (that is, if the byte alignment of **x** is not a  
 1727 multiple of the size of **x**), then the behavior of the atomic region is implementation-defined.

## 1728 Restrictions

- 1729 • All atomic accesses to the storage locations designated by **x** throughout the program are

- 1730 required to have the same type and type parameters.
- 1731 • Storage locations designated by **x** must be less than or equal in size to the largest available
- 1732 native atomic operator width.

## 1733 2.13. Declare Directive

1734 **Summary** A **declare** directive is used in the declaration section of a Fortran subroutine, func-  
 1735 tion, or module, or following a variable declaration in C or C++. It can specify that a *var* is to be  
 1736 allocated in device memory for the duration of the implicit data region of a function, subroutine  
 1737 or program, and specify whether the data values are to be transferred from local memory to device  
 1738 memory upon entry to the implicit data region, and from device memory to local memory upon exit  
 1739 from the implicit data region. These directives create a visible device copy of the *var*.

1740 **Syntax** In C and C++, the syntax of the **declare** directive is:

```
#pragma acc declare clause-list new-line
```

1741 In Fortran the syntax of the **declare** directive is:

```
!$acc declare clause-list
```

1742 where *clause* is one of the following:

```
copy( var-list )
copyin( [readonly:]var-list )
copyout( var-list )
create( var-list )
present( var-list )
deviceptr( var-list )
device_resident( var-list )
link( var-list )
```

1743 The associated region is the implicit region associated with the function, subroutine, or program in  
 1744 which the directive appears. If the directive appears in the declaration section of a Fortran *module*  
 1745 subprogram or in a C or C++ global scope, the associated region is the implicit region for the whole  
 1746 program. The **copy**, **copyin**, **copyout**, **present**, and **deviceptr** data clauses are described  
 1747 in Section 2.7 Data Clauses.

### 1748 Restrictions

- 1749 • A **declare** directive must appear in the same scope as any *var* in any of the data clauses on
- 1750 the directive.
- 1751 • A *var* in a **declare** declare must be a variable or array name, or a Fortran *common block*
- 1752 name between slashes.

- 1753 • A *var* may appear at most once in all the clauses of **declare** directives for a function,  
1754 subroutine, program, or module.
- 1755 • In Fortran, assumed-size dummy arrays may not appear in a **declare** directive.
- 1756 • In Fortran, pointer arrays may be specified, but pointer association is not preserved in device  
1757 memory.
- 1758 • In a Fortran *module* declaration section, only **create**, **copyin**, **device\_resident**, and  
1759 **link** clauses are allowed.
- 1760 • In C or C++ global scope, only **create**, **copyin**, **deviceptr**, **device\_resident** and  
1761 **link** clauses are allowed.
- 1762 • C and C++ *extern* variables may only appear in **create**, **copyin**, **deviceptr**, **device\_resident**  
1763 and **link** clauses on a **declare** directive.
- 1764 • In C and C++, only global and *extern* variables may appear in a **link** clause. In Fortran,  
1765 only *module* variables and *common* block names (enclosed in slashes) may appear in a **link**  
1766 clause.
- 1767 • In C or C++, a **longjmp** call in the region must return to a **setjmp** call within the region.
- 1768 • In C++, an exception thrown in the region must be handled within the region.
- 1769 • See Section 2.17 Fortran Optional Arguments for discussion of Fortran optional dummy ar-  
1770 guments in data clauses, including **device\_resident** clauses.

### 1771 2.13.1. device\_resident clause

1772 **Summary** The **device\_resident** clause specifies that the memory for the named variables  
1773 should be allocated in the current device memory and not in local memory. The host may not be  
1774 able to access variables in a **device\_resident** clause. The accelerator data lifetime of global  
1775 variables or common blocks specified in a **device\_resident** clause is the entire execution of  
1776 the program.

1777 In Fortran, if the variable has the Fortran *allocatable* attribute, the memory for the variable will  
1778 be allocated in and deallocated from the current device memory when the host thread executes  
1779 an **allocate** or **deallocate** statement for that variable, if the current device is a non-shared  
1780 memory device. If the variable has the Fortran *pointer* attribute, it may be allocated or deallocated  
1781 by the host in the current device memory, or may appear on the left hand side of a pointer assignment  
1782 statement, if the right hand side variable itself appears in a **device\_resident** clause.

1783 In Fortran, the argument to a **device\_resident** clause may be a *common block* name enclosed  
1784 in slashes; in this case, all declarations of the common block must have a matching **device\_resident**  
1785 clause. In this case, the *common block* will be statically allocated in device memory, and not  
1786 in local memory. The *common block* will be available to accelerator routines; see Section 2.15  
1787 Procedure Calls in Compute Regions.

1788 In a Fortran *module* declaration section, a *var* in a **device\_resident** clause will be available to  
1789 accelerator subprograms.

1790 In C or C++ global scope, a *var* in a **device\_resident** clause will be available to accelerator  
1791 routines. A C or C++ *extern* variable may appear in a **device\_resident** clause only if the  
1792 actual declaration and all *extern* declarations are also followed by **device\_resident** clauses.

### 1793 2.13.2. create clause

1794 For data in shared memory, no action is taken.

1795 For data not in shared memory, the **create** clause on a **declare** directive behaves as follows,  
1796 for each *var* in *var-list*:

- 1797 • At entry to an implicit data region where the **declare** directive appears:
  - 1798 – If *var* is present, a *present increment* action with the structured reference counter is  
1799 performed. If *var* is a pointer reference, an *attach* action is performed.
  - 1800 – Otherwise, a *create* action with the structured reference counter is performed. If *var* is  
1801 a pointer reference, an *attach* action is performed.
- 1802 • At exit from an implicit data region where the **declare** directive appears:
  - 1803 – If *var* is not present in the current device memory, a runtime error is issued.
  - 1804 – Otherwise, a *present decrement* action with the structured reference counter is per-  
1805 formed. If *var* is a pointer reference, a *detach* action is performed. If both structured  
1806 and dynamic reference counters are zero, a *delete* action is performed.

1807 If the **declare** directive appears in a global context, then the data in *var-list* is statically allocated  
1808 in device memory and the structured reference counter is set to one.

1809 In Fortran, if a variable *var* in *var-list* has the Fortran *allocatable* or *pointer* attribute, then:

- 1810 • An **allocate** statement for *var* will allocate memory in both local memory as well as in the  
1811 current device memory, for a non-shared memory device, and the dynamic reference counter  
1812 will be set to one.
- 1813 • A **deallocate** statement for *var* will deallocate memory from both local memory as well  
1814 as the current device memory, for a non-shared memory device, and the dynamic reference  
1815 counter will be set to zero. If the structured reference counter is not zero, a runtime error is  
1816 issued.

1817 In Fortran, if a variable *var* in *var-list* has the Fortran *pointer* attribute, then it may appear on the  
1818 left hand side of a pointer assignment statement, if the right hand side variable itself appears in a  
1819 **create** clause.

### 1820 2.13.3. link clause

1821 The **link** clause is used for large global host static data that is referenced within an accelerator  
1822 routine and that should have a dynamic data lifetime on the device. The **link** clause specifies that  
1823 only a global link for the named variables should be statically created in accelerator memory. The  
1824 host data structure remains statically allocated and globally available. The device data memory will  
1825 be allocated only when the global variable appears on a data clause for a **data** construct, compute  
1826 construct, or **enter data** directive. The arguments to the **link** clause must be global data. In C  
1827 or C++, the **link** clause must appear at global scope, or the arguments must be *extern* variables.  
1828 In Fortran, the **link** clause must appear in a *module* declaration section, or the arguments must be  
1829 *common block* names enclosed in slashes. A *common block* that is listed in a **link** clause must be  
1830 declared with the same size in all program units where it appears. A **declare link** clause must  
1831 be visible everywhere the global variables or common block variables are explicitly or implicitly

1832 used in a data clause, compute construct, or accelerator routine. The global variable or *common*  
 1833 *block* variables may be used in accelerator routines. The accelerator data lifetime of variables or  
 1834 common blocks specified in a **link** clause is the data region that allocates the variable or common  
 1835 block with a data clause, or from the execution of the **enter data** directive that allocates the data  
 1836 until an **exit data** directive deallocates it or until the end of the program.

## 1837 2.14. Executable Directives

### 1838 2.14.1. Init Directive

1839 **Summary** The **init** directive tells the runtime to initialize the runtime for that device type.  
 1840 This can be used to isolate any initialization cost from the computational cost, when collecting  
 1841 performance statistics. If no device type is specified all devices will be initialized. An **init**  
 1842 directive may be used in place of a call to the **acc\_init** runtime API routine, as described in  
 1843 Section 3.2.7.

1844 **Syntax** In C and C++, the syntax of the **init** directive is:

```
#pragma acc init [clause-list] new-line
```

1845 In Fortran the syntax of the **init** directive is:

```
!$acc init [clause-list]
```

1846 where *clause* is one of the following:

```
device_type ( device-type-list )
device_num ( int-expr )
```

#### 1847 device.type clause

1848 The **device\_type** clause specifies the type of device that is to be initialized in the runtime. If the  
 1849 **device\_type** clause appears, then the *acc-current-device-type-var* for the current thread is set to  
 1850 the argument value. If no **device\_num** clause appears then all devices of this type are initialized.

#### 1851 device.num clause

1852 The **device\_num** clause specifies the device id to be initialized. If the **device\_num** clause  
 1853 appears, then the *acc-current-device-num-var* for the current thread is set to the argument value. If  
 1854 no **device\_type** clause is specified, then the specified device id will be initialized for all available  
 1855 device types.

1856 **Restrictions**

- 1857 • This directive may not be called within a compute region.
- 1858 • If the device type specified is not available, the behavior is implementation-defined; in partic-  
1859 ular, the program may abort.
- 1860 • If the directive is called more than once without an intervening **acc\_shutdown** call or  
1861 **shutdown** directive, with a different value for the device type argument, the behavior is  
1862 implementation-defined.
- 1863 • If some accelerator regions are compiled to only use one device type, using this directive with  
1864 a different device type may produce undefined behavior.

1865 **2.14.2. Shutdown Directive**

1866 **Summary** The **shutdown** directive tells the runtime to shut down the connection to the given  
1867 accelerator, and free any runtime resources. A **shutdown** directive may be used in place of a call  
1868 to the **acc\_shutdown** runtime API routine, as described in Section 3.2.8.

1869 **Syntax** In C and C++, the syntax of the **shutdown** directive is:

```
#pragma acc shutdown [clause-list] new-line
```

1870 In Fortran the syntax of the **shutdown** directive is:

```
!$acc shutdown [clause-list]
```

1871 where *clause* is one of the following:

```
device_type ( device-type-list )  
device_num ( int-expr )
```

1872 **device\_type clause**

1873 The **device\_type** clause specifies the type of device that is to be disconnected from the runtime.  
1874 If no **device\_num** clause appears then all devices of this type are disconnected.

1875 **device\_num clause**

1876 The **device\_num** clause specifies the device id to be disconnected.  
1877 If no clauses appear then all available devices will be disconnected.

1878 **Restrictions**

- 1879 • This directive may not be used during the execution of a compute region.

1880 **2.14.3. Set Directive**

1881 **Summary** The **set** directive provides a means to modify internal control variables using direc-  
 1882 tives. Each form of the **set** directive is functionally equivalent to a matching runtime API routine.

1883 **Syntax** In C and C++, the syntax of the **set** directive is:

```
#pragma acc set [clause-list] new-line
```

1884 In Fortran the syntax of the **set** directive is:

```
!$acc set [clause-list]
```

1885 where *clause* is one of the following

```
default_async ( int-expr )
device_num ( int-expr )
device_type ( device-type-list )
```

1886 **default\_async clause**

1887 The **default\_async** clause specifies the asynchronous queue that should be used if no queue  
 1888 is specified and changes the value of *acc-default-async-var* for the current thread to the argument  
 1889 value. If the value is **acc\_async\_default**, the value of *acc-default-async-var* will revert to  
 1890 the initial value, which is implementation-defined. A **set default\_async** directive is function-  
 1891 ally equivalent to a call to the **acc\_set\_default\_async** runtime API routine, as described in  
 1892 Section 3.2.16.

1893 **device\_num clause**

1894 The **device\_num** clause specifies the device number to set as the default device for accelerator  
 1895 regions and changes the value of *acc-current-device-num-var* for the current thread to the argument  
 1896 value. If the value of **device\_num** argument is negative, the runtime will revert to the default be-  
 1897 havior, which is implementation-defined. A **set device\_num** directive is functionally equivalent  
 1898 to the **acc\_set\_device\_num** runtime API routine, as described in Section 3.2.4.

1899 **device\_type clause**

1900 The **device\_type** clause specifies the device type to set as the default device type for accelerator  
 1901 regions and sets the value of *acc-current-device-type-var* for the current thread to the argument  
 1902 value. If the value of the **device\_type** argument is zero or the clause does not appear, the  
 1903 selected device number will be used for all attached accelerator types. A **set device\_type**  
 1904 directive is functionally equivalent to a call to the **acc\_set\_device\_type** runtime API routine,  
 1905 as described in Section 3.2.2.



1906 **Restrictions**

- 1907 • This directive may not be used within a compute region.
- 1908 • Passing **default\_async** the value of **acc\_async\_noval** has no effect.
- 1909 • Passing **default\_async** the value of **acc\_async\_sync** will cause all asynchronous directives in the default asynchronous queue to become synchronous.
- 1910
- 1911 • Passing **default\_async** the value of **acc\_async\_default** will restore the default asynchronous queue to the initial value, which is implementation-defined.
- 1912
- 1913 • If the value of **device\_num** is larger than the maximum supported value for the given type, the behavior is implementation-defined.
- 1914
- 1915 • At least one clause must appear.
- 1916 • Two instances of the same clause may not appear on the same directive.

1917 **2.14.4. Update Directive**

1918 **Summary** The **update** directive is used during the lifetime of accelerator data to update *vars* in local memory with values from the corresponding data in device memory, or to update *vars* in device memory with values from the corresponding data in local memory.

1919

1920

1921 **Syntax** In C and C++, the syntax of the **update** directive is:

```
#pragma acc update clause-list new-line
```

1922 In Fortran the syntax of the **update** data directive is:

```
!$acc update clause-list
```

1923 where *clause* is one of the following:

```
async [( int-expr )]
wait [( int-expr-list )]
device_type( device-type-list )
if( condition )
if_present
self( var-list )
host( var-list )
device( var-list )
```

1924 Multiple subarrays of the same array may appear in a *var-list* of the same or different clauses on the same directive. The effect of an **update** clause is to copy data from device memory to local memory for **update self**, and from local memory to device memory for **update device**. The updates are done in the order in which they appear on the directive. At least one **self**, **host**, or **device** clause must appear on the directive.

1925

1926

1927

1928

**1929 self clause**

1930 The **self** clause specifies that the *vars* in *var-list* are to be copied from the current device memory  
1931 to local memory for data not in shared memory. For data in shared memory, no action is taken. An  
1932 **update** directive with the **self** clause is equivalent to a call to the **acc\_update\_self** routine,  
1933 described in Section 3.2.25.

**1934 host clause**

1935 The **host** clause is a synonym for the **self** clause.

**1936 device clause**

1937 The **device** clause specifies that the *vars* in *var-list* are to be copied from local memory to the cur-  
1938 rent device memory, for data not in shared memory. For data in shared memory, no action is taken.  
1939 An **update** directive with the **device** clause is equivalent to a call to the **acc\_update\_device**  
1940 routine, described in Section 3.2.24.

**1941 if clause**

1942 The **if** clause is optional; when there is no **if** clause, the implementation will generate code to  
1943 perform the updates unconditionally. When an **if** clause appears, the implementation will generate  
1944 code to conditionally perform the updates only when the *condition* evaluates to nonzero in C or  
1945 C++, or **.true.** in Fortran.

**1946 async clause**

1947 The **async** clause is optional; see Section 2.16 Asynchronous Behavior for more information.

**1948 wait clause**

1949 The **wait** clause is optional; see Section 2.16 Asynchronous Behavior for more information.

**1950 if\_present clause**

1951 When an **if\_present** clause appears on the directive, no action is taken for a *var* which appears  
1952 in *var-list* that is not present in the current device memory. When no **if\_present** clause ap-  
1953 pears, all *vars* in a **device** or **self** clause must be present in the current device memory, and an  
1954 implementation may halt the program with an error message if some data is not present.

**1955 Restrictions**

- 1956 • The **update** directive is executable. It must not appear in place of the statement following  
1957 an *if*, *while*, *do*, *switch*, or *label* in C or C++, or in place of the statement following a logical  
1958 *if* in Fortran.

- 1959 • If no **if\_present** clause appears on the directive, each *var* in *var-list* must be present in  
1960 the current device memory.
- 1961 • Only the **async** and **wait** clauses may follow a **device\_type** clause.
- 1962 • At most one **if** clause may appear. In Fortran, the condition must evaluate to a scalar logical  
1963 value; in C or C++, the condition must evaluate to a scalar integer value.
- 1964 • Noncontiguous subarrays may be specified. It is implementation-specific whether noncon-  
1965 tiguous regions are updated by using one transfer for each contiguous subregion, or whether  
1966 the noncontiguous data is packed, transferred once, and unpacked, or whether one or more  
1967 larger subarrays (no larger than the smallest contiguous region that contains the specified  
1968 subarray) are updated.
- 1969 • In C and C++, a member of a struct or class may be specified, including a subarray of a  
1970 member. Members of a subarray of struct or class type may not be specified.
- 1971 • In C and C++, if a subarray notation is used for a struct member, subarray notation may not  
1972 be used for any parent of that struct member.
- 1973 • In Fortran, members of variables of derived type may be specified, including a subarray of a  
1974 member. Members of subarrays of derived type may not be specified.
- 1975 • In Fortran, if array or subarray notation is used for a derived type member, array or subarray  
1976 notation may not be used for a parent of that derived type member.
- 1977 • See Section 2.17 Fortran Optional Arguments for discussion of Fortran optional arguments in  
1978 **self**, **host**, and **device** clauses.

#### 1979 **2.14.5. Wait Directive**

1980 See Section 2.16 Asynchronous Behavior for more information.

#### 1981 **2.14.6. Enter Data Directive**

1982 See Section 2.6.5 Enter Data and Exit Data Directives for more information.

#### 1983 **2.14.7. Exit Data Directive**

1984 See Section 2.6.5 Enter Data and Exit Data Directives for more information.

### 1985 **2.15. Procedure Calls in Compute Regions**

1986 This section describes how routines are compiled for an accelerator and how procedure calls are  
1987 compiled in compute regions. See Section 2.17 Fortran Optional Arguments for discussion of For-  
1988 tran optional arguments in procedure calls inside compute regions.

### 1989 2.15.1. Routine Directive

1990 **Summary** The **routine** directive is used to tell the compiler to compile a given procedure for  
 1991 an accelerator as well as for the host. In a file or routine with a procedure call, the **routine**  
 1992 directive tells the implementation the attributes of the procedure when called on the accelerator.

1993 **Syntax** In C and C++, the syntax of the **routine** directive is:

```
#pragma acc routine clause-list new-line
#pragma acc routine ( name ) clause-list new-line
```

1994 In C and C++, the **routine** directive without a name may appear immediately before a function  
 1995 definition or just before a function prototype and applies to that immediately following function or  
 1996 prototype. The **routine** directive with a name may appear anywhere that a function prototype  
 1997 is allowed and applies to the function in that scope with that name, but must appear before any  
 1998 definition or use of that function.

1999 In Fortran the syntax of the **routine** directive is:

```
!$acc routine clause-list
!$acc routine ( name ) clause-list
```

2000 In Fortran, the **routine** directive without a name may appear within the specification part of a  
 2001 subroutine or function definition, or within an interface body for a subroutine or function in an  
 2002 interface block, and applies to the containing subroutine or function. The **routine** directive with  
 2003 a name may appear in the specification part of a subroutine, function or module, and applies to the  
 2004 named subroutine or function.

2005 A C or C++ function or Fortran subprogram compiled with the **routine** directive for an accelera-  
 2006 tor is called an *accelerator routine*.

2007 The *clause* is one of the following:

```
gang
worker
vector
seq
bind( name )
bind( string )
device_type( device-type-list )
nohost
```

2008 A **gang**, **worker**, **vector**, or **seq** clause specifies the *level of parallelism* in the routine.

#### 2009 **gang clause**

2010 The **gang** clause specifies that the procedure contains, may contain, or may call another procedure  
 2011 that contains a loop with a **gang** clause. A call to this procedure must appear in code that is

2012 executed in *gang-redundant* mode, and all gangs must execute the call. For instance, a procedure  
2013 with a **routine gang** directive may not be called from within a loop that has a **gang** clause.  
2014 Only one of the **gang**, **worker**, **vector** and **seq** clauses may be specified for each device type.

#### 2015 **worker clause**

2016 The **worker** clause specifies that the procedure contains, may contain, or may call another pro-  
2017 cedure that contains a loop with a **worker** clause, but does not contain nor does it call another  
2018 procedure that contains a loop with the **gang** clause. A loop in this procedure with an **auto** clause  
2019 may be selected by the compiler to execute in **worker** or **vector** mode. A call to this procedure  
2020 must appear in code that is executed in *worker-single* mode, though it may be in *gang-redundant*  
2021 or *gang-partitioned* mode. For instance, a procedure with a **routine worker** directive may be  
2022 called from within a loop that has the **gang** clause, but not from within a loop that has the **worker**  
2023 clause. Only one of the **gang**, **worker**, **vector**, and **seq** clauses may be specified for each  
2024 device type.

#### 2025 **vector clause**

2026 The **vector** clause specifies that the procedure contains, may contain, or may call another pro-  
2027 cedure that contains a loop with the **vector** clause, but does not contain nor does it call another  
2028 procedure that contains a loop with either a **gang** or **worker** clause. A loop in this procedure with  
2029 an **auto** clause may be selected by the compiler to execute in **vector** mode, but not **worker**  
2030 mode. A call to this procedure must appear in code that is executed in *vector-single* mode, though  
2031 it may be in *gang-redundant* or *gang-partitioned* mode, and in *worker-single* or *worker-partitioned*  
2032 mode. For instance, a procedure with a **routine vector** directive may be called from within  
2033 a loop that has the **gang** clause or the **worker** clause, but not from within a loop that has the  
2034 **vector** clause. Only one of the **gang**, **worker**, **vector**, and **seq** clauses may be specified for  
2035 each device type.

#### 2036 **seq clause**

2037 The **seq** clause specifies that the procedure does not contain nor does it call another procedure that  
2038 contains a loop with a **gang**, **worker**, or **vector** clause. A loop in this procedure with an **auto**  
2039 clause will be executed in **seq** mode. A call to this procedure may appear in any mode. Only one  
2040 of the **gang**, **worker**, **vector** and **seq** clauses may be specified for each device type.

#### 2041 **bind clause**

2042 The **bind** clause specifies the name to use when calling the procedure on a device other than the  
2043 host. If the name is specified as an identifier, it is called as if that name were specified in the  
2044 language being compiled. If the name is specified as a string, the string is used for the procedure  
2045 name unmodified. A **bind** clause on a procedure definition behaves as if it had appeared on a  
2046 declaration by changing the name used to call the function on a device other than the host; however,  
2047 the procedure is not compiled for the device with either the original name or the name in the **bind**  
2048 clause.

2049 If there is both a Fortran bind and an acc **bind** clause for a procedure definition then a call on the  
2050 host will call the Fortran bound name and a call on another device will call the name in the **bind**  
2051 clause.

## 2052 **device\_type** clause

2053 The **device\_type** clause is described in Section 2.4 Device-Specific Clauses.

## 2054 **nohost** clause

2055 The **nohost** tells the compiler not to compile a version of this procedure for the host. All calls  
2056 to this procedure must appear within compute regions. If this procedure is called from other pro-  
2057 cedures, those other procedures must also have a matching **routine** directive with the **nohost**  
2058 clause.

## 2059 **Restrictions**

- 2060 • Only the **gang**, **worker**, **vector**, **seq** and **bind** clauses may follow a **device\_type**  
2061 clause.
- 2062 • At least one of the (**gang**, **worker**, **vector**, or **seq**) clauses must appear on the construct.  
2063 If the **device\_type** clause appears on the **routine** directive, a default level of parallelism  
2064 clause must appear before the **device\_type** clause, or a level of parallelism clause must  
2065 be specified following each **device\_type** clause on the directive.
- 2066 • In C and C++, function static variables are not supported in functions to which a **routine**  
2067 directive applies.
- 2068 • In Fortran, variables with the *save* attribute, either explicitly or implicitly, are not supported  
2069 in subprograms to which a **routine** directive applies.
- 2070 • A **bind** clause may not bind to a routine name that has a visible **bind** clause.
- 2071 • If a function or subroutine has a **bind** clause on both the declaration and the definition then  
2072 they both must bind to the same name.

## 2073 **2.15.2. Global Data Access**

2074 C or C++ global, file static, or *extern* variables or array, and Fortran *module* or *common block* vari-  
2075 ables or arrays, that are used in accelerator routines must appear in a declare directive in a **create**,  
2076 **copyin**, **device\_resident** or **link** clause. If the data appears in a **device\_resident**  
2077 clause, the **routine** directive for the procedure must include the **nohost** clause. If the data ap-  
2078 pears in a **link** clause, that data must have an active accelerator data lifetime by virtue of appearing  
2079 in a data clause for a **data** construct, compute construct, or **enter data** directive.

## 2.16. Asynchronous Behavior

This section describes the **async** clause and the behavior of programs that use asynchronous data movement and compute constructs, and asynchronous API routines.

### 2.16.1. **async** clause

The **async** clause may appear on a **parallel**, **kernels**, or **serial** construct, or an **enter data**, **exit data**, **update**, or **wait** directive. In all cases, the **async** clause is optional. When there is no **async** clause on a compute or data construct, the local thread will wait until the compute construct or data operations for the current device are complete before executing any of the code that follows. When there is no **async** clause on a **wait** directive, the local thread will wait until all operations on the appropriate asynchronous activity queues for the current device are complete. When there is an **async** clause, the parallel, kernels, or serial region or data operations may be processed asynchronously while the local thread continues with the code following the construct or directive.

The **async** clause may have a single *async-argument*, where an *async-argument* is a nonnegative scalar integer expression (*int* for C or C++, *integer* for Fortran), or one of the special **async** values defined below. The behavior with a negative *async-argument*, except the special **async** values defined below, is implementation-defined. The value of the *async-argument* may be used in a **wait** directive, **wait** clause, or various runtime routines to test or wait for completion of the operation.

Two special **async** values are defined in the C and Fortran header files and the Fortran **openacc** module. These are negative values, so as not to conflict with a user-specified nonnegative *async-argument*. An **async** clause with the *async-argument* **acc\_async\_noval** will behave the same as if the **async** clause had no argument. An **async** clause with the *async-argument* **acc\_async\_sync** will behave the same as if no **async** clause appeared.

The *async-value* of any operation is the value of the *async-argument*, if it appears, or the value of *acc-default-async-var* if it is **acc\_async\_noval** or if the **async** clause had no value, or **acc\_async\_sync** if no **async** clause appeared. If the current device supports asynchronous operation with one or more device activity queues, the *async-value* is used to select the queue on the current device onto which to enqueue an operation. The properties of the current device and the implementation will determine how many actual activity queues are supported, and how the *async-value* is mapped onto the actual activity queues. Two asynchronous operations with the same current device and the same *async-value* will be enqueued onto the same activity queue, and therefore will be executed on the device in the order they are encountered by the local thread. Two asynchronous operations with different *async-values* may be enqueued onto different activity queues, and therefore may be executed on the device in either order relative to each other. If there are two or more host threads executing and sharing the same device, two asynchronous operations with the same *async-value* will be enqueued on the same activity queue. If the threads are not synchronized with respect to each other, the operations may be enqueued in either order and therefore may execute on the device in either order. Asynchronous operations enqueued to different devices may execute in any order, regardless of the *async-value* used for each.

2119 **2.16.2. wait clause**

2120 The **wait** clause may appear on a **parallel**, **kernels**, or **serial** construct, or an **enter**  
 2121 **data**, **exit data**, or **update** directive. In all cases, the **wait** clause is optional. When there  
 2122 is no **wait** clause, the associated compute or update operations may be enqueued or launched or  
 2123 executed immediately on the device. If there is an argument to the **wait** clause, it must be a list  
 2124 of one or more *async-arguments*. The compute, data or update operation may not be launched or  
 2125 executed until all operations enqueued up to this point by this thread on the associated asynchronous  
 2126 device activity queues have completed. One legal implementation is for the local thread to wait for  
 2127 all the associated asynchronous device activity queues. Another legal implementation is for the  
 2128 local thread to enqueue the compute or update operation in such a way that the operation will  
 2129 not start until the operations enqueued on the associated asynchronous device activity queues have  
 2130 completed.

2131 **2.16.3. Wait Directive**

2132 **Summary** The **wait** directive causes the local thread to wait for completion of asynchronous  
 2133 operations on the current device, such as an accelerator parallel, kernels, or serial region or an  
 2134 **update** directive, or causes one device activity queue to synchronize with one or more other ac-  
 2135 tivity queues on the current device.

2136 **Syntax** In C and C++, the syntax of the **wait** directive is:

```
#pragma acc wait [( int-expr-list )][clause-list] new-line
```

2137 In Fortran the syntax of the **wait** directive is:

```
!$acc wait [( int-expr-list )][clause-list]
```

2138 where *clause* is:

```
async [( int-expr )]
```

2139 The wait argument, if it appears, must be one or more *async-arguments*.

2140 If there is no wait argument and no **async** clause, the local thread will wait until all operations  
 2141 enqueued by this thread on any activity queue on the current device have completed.

2142 If there are one or more *int-expr* expressions and no **async** clause, the local thread will wait until all  
 2143 operations enqueued by this thread on each of the associated device activity queues have completed.

2144 If there are two or more threads executing and sharing the same device, a **wait** directive with no  
 2145 **async** clause will cause the local thread to wait until all of the appropriate asynchronous opera-  
 2146 tions previously enqueued by that thread have completed. To guarantee that operations have been  
 2147 enqueued by other threads requires additional synchronization between those threads. There is no  
 2148 guarantee that all the similar asynchronous operations initiated by other threads will have completed.



2149 If there is an **async** clause, no new operation may be launched or executed on the **async** activ-  
 2150 ity queue on the current device until all operations enqueued up to this point by this thread on the  
 2151 asynchronous activity queues associated with the wait argument have completed. One legal imple-  
 2152 mentation is for the local thread to wait for all the associated asynchronous device activity queues.  
 2153 Another legal implementation is for the thread to enqueue a synchronization operation in such a  
 2154 way that no new operation will start until the operations enqueued on the associated asynchronous  
 2155 device activity queues have completed.

2156 A **wait** directive is functionally equivalent to a call to one of the **acc\_wait**, **acc\_wait\_async**,  
 2157 **acc\_wait\_all** or **acc\_wait\_all\_async** runtime API routines, as described in Sections 3.2.11,  
 2158 3.2.12, 3.2.13 and 3.2.14.

## 2159 2.17. Fortran Optional Arguments

2160 This section refers to the Fortran intrinsic function **PRESENT**. A call to the Fortran intrinsic function  
 2161 **PRESENT(arg)** returns **.true.**, if **arg** is an optional dummy argument and an actual argument  
 2162 for **arg** was present in the argument list of the call site. This should not be confused with the  
 2163 OpenACC **present** data clause.

2164 The appearance of a Fortran optional argument **arg** as a *var* in any of the following clauses has no  
 2165 effect at runtime if **PRESENT(arg)** is **.false.**:

- 2166 • in data clauses on compute and **data** constructs;
- 2167 • in data clauses on **enter data** and **exit data** directives;
- 2168 • in data and **device\_resident** clauses on **declare** directives;
- 2169 • in **use\_device** clauses on **host\_data** directives;
- 2170 • in **self**, **host**, and **device** clauses on **update** directives.

2171 The appearance of a Fortran optional argument **arg** in the following situations may result in unde-  
 2172 fined behavior if **PRESENT(arg)** is **.false.** when the associated construct is executed:

- 2173 • as a *var* in **private**, **firstprivate**, and **reduction** clauses;
- 2174 • as a *var* in **cache** directives;
- 2175 • as part of an expression in any clause or directive.

2176 A call to the Fortran intrinsic function **PRESENT** behaves the same way in a compute construct or  
 2177 an accelerator routine as on the host. The function call **PRESENT(arg)** must return the same value  
 2178 in a compute construct as **PRESENT(arg)** would outside of the compute construct. If a Fortran  
 2179 optional argument **arg** appears as an actual argument in a procedure call in a compute construct  
 2180 or an accelerator routine, and the associated dummy argument **subarg** also has the **optional**  
 2181 attribute, then **PRESENT(subarg)** returns the same value as **PRESENT(subarg)** would when  
 2182 executed on the host.



## 2183 3. Runtime Library

2184 This chapter describes the OpenACC runtime library routines that are available for use by program-  
2185 mers. Use of these routines may limit portability to systems that do not support the OpenACC API.  
2186 Conditional compilation using the `_OPENACC` preprocessor variable may preserve portability.

2187 This chapter has two sections:

- 2188 • Runtime library definitions
- 2189 • Runtime library routines

2190 There are four categories of runtime routines:

- 2191 • Device management routines, to get the number of devices, set the current device, and so on.
- 2192 • Asynchronous queue management, to synchronize until all activities on an async queue are  
2193 complete, for instance.
- 2194 • Device test routine, to test whether this statement is executing on the device or not.
- 2195 • Data and memory management, to manage memory allocation or copy data between memo-  
2196 ries.

### 2197 3.1. Runtime Library Definitions

2198 In C and C++, prototypes for the runtime library routines described in this chapter are provided in  
2199 a header file named `openacc.h`. All the library routines are *extern* functions with “C” linkage.  
2200 This file defines:

- 2201 • The prototypes of all routines in the chapter.
- 2202 • Any datatypes used in those prototypes, including an enumeration type to describe the sup-  
2203 ported device types.
- 2204 • The values of `acc_async_noval`, `acc_async_sync`, and `acc_async_default`.

2205 In Fortran, interface declarations are provided in a Fortran module named `openacc`. The `openacc`  
2206 module defines:

- 2207 • The integer parameter `openacc_version` with a value `yyyymm` where `yyyy` and `mm` are the  
2208 year and month designations of the version of the Accelerator programming model supported.  
2209 This value matches the value of the preprocessor variable `_OPENACC`.
- 2210 • Interfaces for all routines in the chapter.
- 2211 • Integer parameters to define integer kinds for arguments to and return values for those rou-  
2212 tines.

- 2213 • Integer parameters to describe the supported device types.
- 2214 • Integer parameters to define the values of `acc_async_noval`, `acc_async_sync`, and
- 2215 `acc_async_default`.

2216 Many of the routines accept or return a value corresponding to the type of device. In C and C++, the  
 2217 datatype used for device type values is `acc_device_t`; in Fortran, the corresponding datatype  
 2218 is `integer(kind=acc_device_kind)`. The possible values for device type are implemen-  
 2219 tation specific, and are defined in the C or C++ include file `openacc.h` and the Fortran module  
 2220 `openacc`. Four values are always supported: `acc_device_none`, `acc_device_default`,  
 2221 `acc_device_host` and `acc_device_not_host`. For other values, look at the appropriate  
 2222 files included with the implementation, or read the documentation for the implementation. The  
 2223 value `acc_device_default` will never be returned by any function; its use as an argument will  
 2224 tell the runtime library to use the default device type for that implementation.

## 2225 3.2. Runtime Library Routines

2226 In this section, for the C and C++ prototypes, pointers are typed `h_void*` or `d_void*` to desig-  
 2227 nate a host memory address or device memory address, when these calls are executed on the host,  
 2228 as if the following definitions were included:

```
#define h_void void
#define d_void void
```

2229 Except for `acc_on_device`, these routines are only available on the host.

### 2230 3.2.1. `acc_get_num_devices`

2231 **Summary** The `acc_get_num_devices` routine returns the number of devices of the given  
 2232 type available.

#### 2233 **Format**

C or C++:

```
int acc_get_num_devices( acc_device_t );
```

Fortran:

```
integer function acc_get_num_devices( devicetype )
  integer(acc_device_kind) :: devicetype
```

2234 **Description** The `acc_get_num_devices` routine returns the number of devices of the given  
 2235 type available. The argument tells what kind of device to count.

#### 2236 **Restrictions**

- 2237 • This routine may not be called within a compute region.

### 2238 3.2.2. `acc_set_device_type`

2239 **Summary** The `acc_set_device_type` routine tells the runtime which type of device to use  
2240 when executing a compute region and sets the value of `acc-current-device-type-var`. This is useful  
2241 when the implementation allows the program to be compiled to use more than one type of device.

#### 2242 **Format**

C or C++:

```
void acc_set_device_type( acc_device_t );
```

Fortran:

```
subroutine acc_set_device_type( devicetype )  
integer(acc_device_kind) :: devicetype
```

2243 **Description** The `acc_set_device_type` routine tells the runtime which type of device to  
2244 use among those available and sets the value of `acc-current-device-type-var` for the current thread.  
2245 A call to `acc_set_device_type` is functionally equivalent to a `set device_type` directive  
2246 with the matching device type argument, as described in Section 2.14.3.

#### 2247 **Restrictions**

- 2248 • If the device type specified is not available, the behavior is implementation-defined; in partic-  
2249 ular, the program may abort.
- 2250 • If some compute regions are compiled to only use one device type, calling this routine with a  
2251 different device type may produce undefined behavior.

### 2252 3.2.3. `acc_get_device_type`

2253 **Summary** The `acc_get_device_type` routine returns the value of `acc-current-device-type-`  
2254 `var`, which is the device type of the current device. This is useful when the implementation allows  
2255 the program to be compiled to use more than one type of device.

#### 2256 **Format**

C or C++:

```
acc_device_t acc_get_device_type( void );
```

Fortran:

```
function acc_get_device_type()  
integer(acc_device_kind) :: acc_get_device_type
```

2257 **Description** The `acc_get_device_type` routine returns the value of *acc-current-device-*  
 2258 *type-var* for the current thread to tell the program what type of device will be used to run the next  
 2259 compute region, if one has been selected. The device type may have been selected by the program  
 2260 with an `acc_set_device_type` call, with an environment variable, or by the default behavior  
 2261 of the program.

#### 2262 **Restrictions**

- 2263 • If the device type has not yet been selected, the value `acc_device_none` may be returned.

### 2264 **3.2.4. acc\_set\_device\_num**

2265 **Summary** The `acc_set_device_num` routine tells the runtime which device to use and sets  
 2266 the value of *acc-current-device-num-var*.

#### 2267 **Format**

C or C++:

```
void acc_set_device_num( int, acc_device_t );
```

Fortran:

```
subroutine acc_set_device_num( devicenum, devicetype )
  integer :: devicenum
  integer(acc_device_kind) :: devicetype
```

2268 **Description** The `acc_set_device_num` routine tells the runtime which device to use among  
 2269 those available of the given type for compute or data regions in the current thread and sets the value  
 2270 of *acc-current-device-num-var*. If the value of `devicenum` is negative, the runtime will revert to  
 2271 its default behavior, which is implementation-defined. If the value of the second argument is zero,  
 2272 the selected device number will be used for all device types. A call to `acc_set_device_num`  
 2273 is functionally equivalent to a `set device_num` directive with the matching device number argu-  
 2274 ment, as described in Section 2.14.3.

#### 2275 **Restrictions**

- 2276 • If the value of `devicenum` is greater than or equal to the value returned by `acc_get_num_devices`  
 2277 for that device type, the behavior is implementation-defined.
- 2278 • Calling `acc_set_device_num` implies a call to `acc_set_device_type` with that  
 2279 device type argument.

### 2280 **3.2.5. acc\_get\_device\_num**

2281 **Summary** The `acc_get_device_num` routine returns the value of *acc-current-device-num-*  
 2282 *var* for the current thread.

2283 **Format**

C or C++:

```
int acc_get_device_num( acc_device_t );
```

Fortran:

```
integer function acc_get_device_num( devicetype )
  integer(acc_device_kind) :: devicetype
```

2284 **Description** The `acc_get_device_num` routine returns the value of *acc-current-device-num-*  
2285 *var* for the current thread.

2286 **3.2.6. acc\_get\_property**

2287 **Summary** The `acc_get_property` and `acc_get_property_string` routines return  
2288 the value of a *device-property* for the specified device.

2289 **Format**

C or C++:

```
size_t acc_get_property( int devicenum,
  acc_device_t devicetype, acc_device_property_t property );
const char* acc_get_property_string( int devicenum,
  acc_device_t devicetype, acc_device_property_t property );
```

Fortran:

```
function acc_get_property( devicenum, devicetype, property )
subroutine acc_get_property_string( devicenum, devicetype,
  property, string )
  integer, value :: devicenum
  integer(acc_device_kind), value :: devicetype
  integer(acc_device_property), value :: property
  integer(acc_device_property) :: acc_get_property
  character(*) :: string
```

2290 **Description** The `acc_get_property` and `acc_get_property_string` routines returns  
2291 the value of the specified *property*. `devicenum` and `devicetype` specify the device being  
2292 queried. If `devicetype` has the value `acc_device_current`, then `devicenum` is ignored  
2293 and the value of the property for the current device is returned. `property` is an enumeration  
2294 constant, defined in `openacc.h`, for C or C++, or an integer parameter, defined in the `openacc`  
2295 module, for Fortran. Integer-valued properties are returned by `acc_get_property`, and string-  
2296 valued properties are returned by `acc_get_property_string`. In Fortran, `acc_get_property_string`  
2297 returns the result into the `character` variable passed as the last argument.

2298 The supported values of `property` are given in the following table.

<i>property</i>	<i>return type</i>	<i>return value</i>
<b>acc_property_memory</b>	<i>integer</i>	size of device memory in bytes
<b>acc_property_free_memory</b>	<i>integer</i>	free device memory in bytes
<b>acc_property_shared_memory_support</b>	<i>integer</i>	nonzero if the specified device supports sharing memory with the local thread
<b>acc_property_name</b>	<i>string</i>	device name
<b>acc_property_vendor</b>	<i>string</i>	device vendor
<b>acc_property_driver</b>	<i>string</i>	device driver version

2300 An implementation may support additional properties for some devices.

### 2301 Restrictions

- 2302 • These routines may not be called within an compute region.
- 2303 • If the value of **property** is not one of the known values for that query routine, or that  
2304 property has no value for the specified device, **acc\_get\_property** will return 0 and  
2305 **acc\_get\_property\_string** will return NULL (in C or C++) or an blank string (in  
2306 Fortran).

### 2307 3.2.7. acc\_init

2308 **Summary** The **acc\_init** routine tells the runtime to initialize the runtime for that device type.  
2309 This can be used to isolate any initialization cost from the computational cost, when collecting  
2310 performance statistics.

### 2311 Format

C or C++:

```
void acc_init( acc_device_t );
```

Fortran:

```
subroutine acc_init( devicetype )
  integer(acc_device_kind) :: devicetype
```

2312 **Description** The **acc\_init** routine also implicitly calls **acc\_set\_device\_type**. A call to  
2313 **acc\_init** is functionally equivalent to a **init** directive with the matching device type argument,  
2314 as described in Section 2.14.1.

### 2315 Restrictions

- 2316 • This routine may not be called within a compute region.
- 2317 • If the device type specified is not available, the behavior is implementation-defined; in partic-  
2318 ular, the program may abort.



- 2319       • If the routine is called more than once without an intervening **acc\_shutdown** call, with a  
2320       different value for the device type argument, the behavior is implementation-defined.
- 2321       • If some accelerator regions are compiled to only use one device type, calling this routine with  
2322       a different device type may produce undefined behavior.

### 2323 3.2.8. **acc\_shutdown**

2324 **Summary** The **acc\_shutdown** routine tells the runtime to shut down any connection to de-  
2325 vices of the given device type, and free up any runtime resources. A call to **acc\_shutdown**  
2326 is functionally equivalent to a **shutdown** directive with the matching device type argument, as  
2327 described in Section 2.14.2.

#### 2328 **Format**

C or C++:

```
void acc_shutdown( acc_device_t );
```

Fortran:

```
subroutine acc_shutdown( devicetype )  
integer(acc_device_kind) :: devicetype
```

2329 **Description** The **acc\_shutdown** routine disconnects the program from any device of the spec-  
2330 ified device type. Any data that is present in the memory of any such device is immediately deallo-  
2331 cated.

#### 2332 **Restrictions**

- 2333       • This routine may not be called during execution of a compute region.
- 2334       • If the program attempts to execute a compute region on a device or to access any data in  
2335       the memory of a device after a call to **acc\_shutdown** for that device type, the behavior is  
2336       undefined.
- 2337       • If the program attempts to shut down the **acc\_device\_host** device type, the behavior is  
2338       undefined.

### 2339 3.2.9. **acc\_async\_test**

2340 **Summary** The **acc\_async\_test** routine tests for completion of all associated asynchronous  
2341 operations on the current device.

#### 2342 **Format**

C or C++:

```
int acc_async_test( int );
```

Fortran:

```
logical function acc_async_test( arg )  
integer(acc_handle_kind) :: arg
```

2343 **Description** The argument must be an *async-argument* as defined in Section 2.16.1 *async* clause.  
2344 If that value did not appear in any **async** clauses, or if it did appear in one or more **async** clauses  
2345 and all such asynchronous operations have completed on the current device, the **acc\_async\_test**  
2346 routine will return with a nonzero value in C and C++, or **.true.** in Fortran. If some such asyn-  
2347 chronous operations have not completed, the **acc\_async\_test** routine will return with a zero  
2348 value in C and C++, or **.false.** in Fortran. If two or more threads share the same accelerator, the  
2349 **acc\_async\_test** routine will return with a nonzero value or **.true.** only if all matching asyn-  
2350 chronous operations initiated by this thread have completed; there is no guarantee that all matching  
2351 asynchronous operations initiated by other threads have completed.

### 2352 3.2.10. acc\_async\_test\_all

2353 **Summary** The **acc\_async\_test\_all** routine tests for completion of all asynchronous op-  
2354 erations.

#### 2355 Format

C or C++:

```
int acc_async_test_all( );
```

Fortran:

```
logical function acc_async_test_all( )
```

2356 **Description** If all outstanding asynchronous operations have completed, the **acc\_async\_test\_all**  
2357 routine will return with a nonzero value in C and C++, or **.true.** in Fortran. If some asynchronous  
2358 operations have not completed, the **acc\_async\_test\_all** routine will return with a zero value  
2359 in C and C++, or **.false.** in Fortran. If two or more threads share the same accelerator, the  
2360 **acc\_async\_test\_all** routine will return with a nonzero value or **.true.** only if all outstand-  
2361 ing asynchronous operations initiated by this thread have completed; there is no guarantee that all  
2362 asynchronous operations initiated by other threads have completed.

### 2363 3.2.11. acc\_wait

2364 **Summary** The **acc\_wait** routine waits for completion of all associated asynchronous opera-  
2365 tions on the current device.

#### 2366 Format

C or C++:

```
void acc_wait( int );
```

Fortran:

```
subroutine acc_wait( arg )
  integer(acc_handle_kind) :: arg
```

2367 **Description** The argument must be an *async-argument* as defined in Section 2.16.1 *async* clause.  
2368 If that value appeared in one or more **async** clauses, the **acc\_wait** routine will not return until  
2369 the latest such asynchronous operation has completed on the current device. If two or more threads  
2370 share the same accelerator, the **acc\_wait** routine will return only if all matching asynchronous  
2371 operations initiated by this thread have completed; there is no guarantee that all matching asyn-  
2372 chronous operations initiated by other threads have completed. For compatibility with version 1.0,  
2373 this routine may also be spelled **acc\_async\_wait**. A call to **acc\_wait** is functionally equiv-  
2374 alent to a **wait** directive with a matching wait argument and no **async** clause, as described in  
2375 Section 2.16.3.

### 2376 3.2.12. acc\_wait\_async

2377 **Summary** The **acc\_wait\_async** routine enqueues a wait operation on one *async* queue of  
2378 the current device for the operations previously enqueued on another *async* queue.

#### 2379 Format

C or C++:

```
void acc_wait_async( int, int );
```

Fortran:

```
subroutine acc_wait_async( arg, async )
  integer(acc_handle_kind) :: arg, async
```

2380 **Description** The arguments must be *async-arguments*, as defined in Section 2.16.1 *async* clause.  
2381 The routine will enqueue a wait operation on the appropriate device queue associated with the  
2382 second argument, which will wait for operations enqueued on the device queue associated with  
2383 the first argument. See Section 2.16 Asynchronous Behavior for more information. A call to  
2384 **acc\_wait\_async** is functionally equivalent to a **wait** directive with a matching wait argument  
2385 and a matching **async** argument, as described in Section 2.16.3.

### 2386 3.2.13. acc\_wait\_all

2387 **Summary** The **acc\_wait\_all** routine waits for completion of all asynchronous operations.

#### 2388 Format

C or C++:

```
void acc_wait_all( );
```

Fortran:

```
subroutine acc_wait_all( )
```

2389 **Description** The `acc_wait_all` routine will not return until the all asynchronous operations  
 2390 have completed. If two or more threads share the same accelerator, the `acc_wait_all` routine  
 2391 will return only if all asynchronous operations initiated by this thread have completed; there is no  
 2392 guarantee that all asynchronous operations initiated by other threads have completed. For com-  
 2393 patibility with version 1.0, this routine may also be spelled `acc_async_wait_all`. A call to  
 2394 `acc_wait_all` is functionally equivalent to a `wait` directive with no wait argument list and no  
 2395 `async` argument, as described in Section 2.16.3.

### 2396 3.2.14. `acc_wait_all_async`

2397 **Summary** The `acc_wait_all_async` routine enqueues wait operations on one async queue  
 2398 for the operations previously enqueued on all other async queues.

#### 2399 Format

C or C++:

```
void acc_wait_all_async( int );
```

Fortran:

```
subroutine acc_wait_all_async( async )  
integer(acc_handle_kind) :: async
```

2400 **Description** The argument must be an *async-argument* as defined in Section 2.16.1 *async* clause.  
 2401 The routine will enqueue a wait operation on the appropriate device queue for each other device  
 2402 queue. See Section 2.16 Asynchronous Behavior for more information. A call to `acc_wait_all_async`  
 2403 is functionally equivalent to a `wait` directive with no wait argument list and a matching `async`  
 2404 argument, as described in Section 2.16.3.

### 2405 3.2.15. `acc_get_default_async`

2406 **Summary** The `acc_get_default_async` routine returns the value of *acc-default-async-*  
 2407 *var* for the current thread.

#### 2408 Format

C or C++:

```
int acc_get_default_async( void );
```

Fortran:

```
function acc_get_default_async( )  
integer(acc_handle_kind) :: acc_get_default_async
```

2409 **Description** The `acc_get_default_async` routine returns the value of *acc-default-async-*  
2410 *var* for the current thread, which is the asynchronous queue used when an **async** clause appears  
2411 without an *async-argument* or with the value `acc_async_noval`.

### 2412 3.2.16. `acc_set_default_async`

2413 **Summary** The `acc_set_default_async` routine tells the runtime which asynchronous queue  
2414 to use when no other queue is specified.

#### 2415 **Format**

C or C++:

```
void acc_set_default_async( int async );
```

Fortran:

```
subroutine acc_set_default_async( async )  
integer(acc_handle_kind) :: async
```

2416 **Description** The `acc_set_default_async` routine tells the runtime to place any directives  
2417 with an **async** clause that does not have an *async-argument* or with the special `acc_async_noval`  
2418 value into the specified asynchronous activity queue instead of the default asynchronous activity  
2419 queue for that device by setting the value of *acc-default-async-var* for the current thread. The spe-  
2420 cial argument `acc_async_default` will reset the default asynchronous activity queue to the  
2421 initial value, which is implementation-defined. A call to `acc_set_default_async` is func-  
2422 tionally equivalent to a `set default_async` directive with a matching argument in *int-expr*, as  
2423 described in Section 2.14.3.

### 2424 3.2.17. `acc_on_device`

2425 **Summary** The `acc_on_device` routine tells the program whether it is executing on a partic-  
2426 ular device.

#### 2427 **Format**

C or C++:

```
int acc_on_device( acc_device_t );
```

Fortran:

```
logical function acc_on_device( devicetype )  
integer(acc_device_kind) :: devicetype
```

2428 **Description** The `acc_on_device` routine may be used to execute different paths depend-  
2429 ing on whether the code is running on the host or on some accelerator. If the `acc_on_device`  
2430 routine has a compile-time constant argument, it evaluates at compile time to a constant. The ar-  
2431 gument must be one of the defined accelerator types. If the argument is `acc_device_host`,  
2432 then outside of a compute region or accelerator routine, or in a compute region or accelerator rou-  
2433 tine that is executed on the host CPU, this routine will evaluate to nonzero for C or C++, and  
2434 `.true.` for Fortran; otherwise, it will evaluate to zero for C or C++, and `.false.` for Fortran.  
2435 If the argument is `acc_device_not_host`, the result is the negation of the result with argu-  
2436 ment `acc_device_host`. If the argument is an accelerator device type, then in a compute region  
2437 or routine that is executed on a device of that type, this routine will evaluate to nonzero for C or  
2438 C++, and `.true.` for Fortran; otherwise, it will evaluate to zero for C or C++, and `.false.` for  
2439 Fortran. The result with argument `acc_device_default` is undefined.

### 2440 3.2.18. `acc_malloc`

2441 **Summary** The `acc_malloc` routine allocates space in the current device memory.

#### 2442 **Format**

C or C++:

```
d_void* acc_malloc( size_t );
```

2443 **Description** The `acc_malloc` routine may be used to allocate space in the current device  
2444 memory. Pointers assigned from this function may be used in `deviceptr` clauses to tell the  
2445 compiler that the pointer target is resident on the device. In case of an error, `acc_malloc` returns  
2446 a NULL pointer.

### 2447 3.2.19. `acc_free`

2448 **Summary** The `acc_free` routine frees memory on the current device.

#### 2449 **Format**

C or C++:

```
void acc_free( d_void* );
```

2450 **Description** The `acc_free` routine will free previously allocated space in the current device  
2451 memory; the argument should be a pointer value that was returned by a call to `acc_malloc`. If  
2452 the argument is a NULL pointer, no operation is performed.

### 2453 3.2.20. `acc_copyin`

2454 **Summary** The `acc_copyin` routines test to see if the argument is in shared memory or already  
2455 present in the current device memory; if not, they allocate space in the current device memory to

2456 correspond to the specified local memory, and copy the data to that device memory.

## 2457 **Format**

C or C++:

```
d_void* acc_copyin( h_void*, size_t );
void acc_copyin_async( h_void*, size_t, int );
```

Fortran:

```
subroutine acc_copyin( a )
subroutine acc_copyin( a, len )
subroutine acc_copyin_async( a, async )
subroutine acc_copyin_async( a, len, async )
  type(*), dimension(..) :: a
  integer :: len
  integer(acc_handle_kind) :: async
```

2458 **Description** The **acc\_copyin** routines are equivalent to the **enter data** directive with a  
 2459 **copyin** clause, as described in Section 2.7.6. In C, the arguments are a pointer to the data and  
 2460 length in bytes; the synchronous function returns a pointer to the allocated device memory, as with  
 2461 **acc\_malloc**. In Fortran, two forms are supported. In the first, the argument is a contiguous array  
 2462 section of intrinsic type. In the second, the first argument is a variable or array element and the  
 2463 second is the length in bytes.

2464 The behavior of the **acc\_copyin** routines is:

- 2465 • If the data is in shared memory, no action is taken. The C **acc\_copyin** returns the incoming  
 2466 pointer.
- 2467 • If the data is present in the current device memory, a *present increment* action with the dy-  
 2468 namic reference counter is performed. The C **acc\_copyin** returns a pointer to the existing  
 2469 device memory.
- 2470 • Otherwise, a *copyin* action with the appropriate reference counter is performed. The C  
 2471 **acc\_copyin** returns the device address of the newly allocated memory.

2472 This data may be accessed using the **present** data clause. Pointers assigned from the C **acc\_copyin**  
 2473 function may be used in **deviceptr** clauses to tell the compiler that the pointer target is resident  
 2474 on the device.

2475 The **\_async** versions of this function will perform any data transfers asynchronously on the **async**  
 2476 queue associated with the value passed in as the **async** argument. The function may return be-  
 2477 fore the data has been transferred; see Section 2.16 Asynchronous Behavior for more details. The  
 2478 synchronous versions will not return until the data has been completely transferred.

2479 For compatibility with OpenACC 2.0, **acc\_present\_or\_copyin** and **acc\_pcopyin** are al-  
 2480 ternate names for **acc\_copyin**.

2481 **3.2.21. acc\_create**

2482 **Summary** The **acc\_create** routines test to see if the argument is in shared memory or already  
 2483 present in the current device memory; if not, they allocate space in the current device memory to  
 2484 correspond to the specified local memory.

2485 **Format**

C or C++:

```
d_void* acc_create( h_void*, size_t );
void acc_create_async( h_void*, size_t, int async );
```

Fortran:

```
subroutine acc_create( a )
subroutine acc_create( a, len )
subroutine acc_create_async( a, async )
subroutine acc_create_async( a, len, async )
  type(*), dimension(..) :: a
  integer :: len
  integer(acc_handle_kind) :: async
```

2486 **Description** The **acc\_create** routines are equivalent to the **enter data** directive with a  
 2487 **create** clause, as described in Section 2.7.8. In C, the arguments are a pointer to the data and  
 2488 length in bytes; the synchronous function returns a pointer to the allocated device memory, as with  
 2489 **acc\_malloc**. In Fortran, two forms are supported. In the first, the argument is a contiguous array  
 2490 section of intrinsic type. In the second, the first argument is a variable or array element and the  
 2491 second is the length in bytes.

2492 The behavior of the **acc\_create** routines is:

- 2493 • If the data is in shared memory, no action is taken. The C **acc\_create** returns the incoming  
 2494 pointer.
- 2495 • If the data is present in the current device memory, a *present increment* action with the dy-  
 2496 namic reference counter is performed. The C **acc\_create** returns a pointer to the existing  
 2497 device memory.
- 2498 • Otherwise, a *create* action with the appropriate reference counter is performed. The C **acc\_create**  
 2499 returns the device address of the newly allocated memory.

2500 This data may be accessed using the **present** data clause. Pointers assigned from the C **acc\_copyin**  
 2501 function may be used in **deviceptr** clauses to tell the compiler that the pointer target is resident  
 2502 on the device.

2503 The **\_async** versions of these function may perform the data allocation asynchronously on the  
 2504 async queue associated with the value passed in as the **async** argument. The synchronous versions  
 2505 will not return until the data has been allocated.

2506 For compatibility with OpenACC 2.0, **acc\_present\_or\_create** and **acc\_pcreate** are al-  
 2507 ternate names for **acc\_create**.



2508 **3.2.22. acc\_copyout**

2509 **Summary** The `acc_copyout` routines test to see if the argument is in shared memory; if not,  
 2510 the argument must be present in the current device memory, and the routines copy data from device  
 2511 memory to the corresponding local memory, then deallocate that space from the device memory.

2512 **Format**

C or C++:

```
void acc_copyout( h_void*, size_t );
void acc_copyout_async( h_void*, size_t, int async );
void acc_copyout_finalize( h_void*, size_t );
void acc_copyout_finalize_async( h_void*, size_t, int async );
```

Fortran:

```
subroutine acc_copyout( a )
subroutine acc_copyout( a, len )
subroutine acc_copyout_async( a, async )
subroutine acc_copyout_async( a, len, async )
subroutine acc_copyout_finalize( a )
subroutine acc_copyout_finalize( a, len )
subroutine acc_copyout_finalize_async( a, async )
subroutine acc_copyout_finalize_async( a, len, async )
type(*), dimension(..) :: a
integer :: len
integer(acc_handle_kind) :: async
```

2513 **Description** The `acc_copyout` routines are equivalent to the `exit data` directive with a  
 2514 `copyout` clause, and the `acc_copyout_finalize` routines are equivalent to the `exit data`  
 2515 directive with both `copyout` and `finalize` clauses, as described in Section 2.7.7. In C, the  
 2516 arguments are a pointer to the data and length in bytes. In Fortran, two forms are supported. In the  
 2517 first, the argument is a contiguous array section of intrinsic type. In the second, the first argument  
 2518 is a variable or array element and the second is the length in bytes.

2519 The behavior of the `acc_copyout` routines is:

- 2520 • If the data is in shared memory, no action is taken.
- 2521 • Otherwise, if the data is not present in the current device memory, a runtime error is issued.
- 2522 • Otherwise, a *present decrement* action with the dynamic reference counter is performed (`acc_copyout`),  
 2523 or the dynamic reference counter is set to zero (`acc_copyout_finalize`). If both ref-  
 2524 erence counters are then zero, a *copyout* action is performed.

2525 The `_async` versions of these functions will perform any associated data transfers asynchronously  
 2526 on the `async` queue associated with the value passed in as the `async` argument. The function may  
 2527 return before the data has been transferred or deallocated; see Section 2.16 Asynchronous Behavior  
 2528 for more details. The synchronous versions will not return until the data has been completely trans-  
 2529 ferred. Even if the data has not been transferred or deallocated before the function returns, the data  
 2530 will be treated as not present in the current device memory.

2531 **3.2.23. acc\_delete**

2532 **Summary** The **acc\_delete** routines test to see if the argument is in shared memory; if not,  
 2533 the argument must be present in the current device memory, and the routines deallocate that space  
 2534 from the device memory.

2535 **Format**

C or C++:

```
void acc_delete( h_void*, size_t );
void acc_delete_async( h_void*, size_t, int async );
void acc_delete_finalize( h_void*, size_t );
void acc_delete_finalize_async( h_void*, size_t, int async );
```

Fortran:

```
subroutine acc_delete( a )
subroutine acc_delete( a, len )
subroutine acc_delete_async( a, async )
subroutine acc_delete_async( a, len, async )
subroutine acc_delete_finalize( a )
subroutine acc_delete_finalize( a, len )
subroutine acc_delete_finalize_async( a, async )
subroutine acc_delete_finalize_async( a, len, async )
  type(*), dimension(..) :: a
  integer :: len
  integer(acc_handle_kind) :: async
```

2536 **Description** The **acc\_delete** routines are equivalent to the **exit data** directive with a  
 2537 **delete** clause,

2538 and the **acc\_delete\_finalize** routines are equivalent to the **exit data** directive with both  
 2539 **delete** clause and **finalize** clauses, as described in Section 2.7.10. The arguments are as for  
 2540 **acc\_copyout**.

2541 The behavior of the **acc\_delete** routines is:

- 2542 • If the data is in shared memory, no action is taken.
- 2543 • Otherwise, if the data is not present in the current device memory, a runtime error is issued.
- 2544 • Otherwise, a *present decrement* action with the dynamic reference counter is performed (**acc\_delete**),  
 2545 or the dynamic reference counter is set to zero (**acc\_delete\_finalize**). If both refer-  
 2546 ence counters are then zero, a *delete* action is performed.

2547 The **\_async** versions of these function may perform the data deallocation asynchronously on the  
 2548 **async** queue associated with the value passed in as the **async** argument. The synchronous versions  
 2549 will not return until the data has been deallocated. Even if the data has not been deallocated before  
 2550 the function returns, the data will be treated as not present in the current device memory.

2551 **3.2.24. acc\_update\_device**

2552 **Summary** The `acc_update_device` routines test to see if the argument is in shared memory;  
 2553 if not, the argument must be present in the current device memory, and the routines update the data  
 2554 in device memory from the corresponding local memory.

2555 **Format**

C or C++:

```
void acc_update_device( h_void*, size_t );
void acc_update_device_async( h_void*, size_t, int async );
```

Fortran:

```
subroutine acc_update_device( a )
subroutine acc_update_device( a, len )
subroutine acc_update_device( a, async )
subroutine acc_update_device( a, len, async )
  type(*), dimension(..) :: a
  integer :: len
  integer(acc_handle_kind) :: async
```

2556 **Description** The `acc_update_device` routine is equivalent to the `update` directive with a  
 2557 `device` clause, as described in Section 2.14.4. In C, the arguments are a pointer to the data and  
 2558 length in bytes. In Fortran, two forms are supported. In the first, the argument is a contiguous array  
 2559 section of intrinsic type. In the second, the first argument is a variable or array element and the  
 2560 second is the length in bytes. For data not in shared memory, the data in the local memory is copied  
 2561 to the corresponding device memory. It is a runtime error to call this routine if the data is not present  
 2562 in the current device memory.

2563 The `_async` versions of this function will perform the data transfers asynchronously on the `async`  
 2564 queue associated with the value passed in as the `async` argument. The function may return be-  
 2565 fore the data has been transferred; see Section 2.16 Asynchronous Behavior for more details. The  
 2566 synchronous versions will not return until the data has been completely transferred.

2567 **3.2.25. acc\_update\_self**

2568 **Summary** The `acc_update_self` routines test to see if the argument is in shared memory;  
 2569 if not, the argument must be present in the current device memory, and the routines update the data  
 2570 in local memory from the corresponding device memory.

2571 **Format**

C or C++:

```
void acc_update_self( h_void*, size_t );
void acc_update_self_async( h_void*, size_t, int async );
```

Fortran:

```

subroutine acc_update_self( a )
subroutine acc_update_self( a, len )
subroutine acc_update_self_async( a, async )
subroutine acc_update_self_async( a, len, async )
  type(*) , dimension(..) :: a
  integer :: len
  integer(acc_handle_kind) :: async

```

2572 **Description** The **acc\_update\_self** routine is equivalent to the **update** directive with a  
 2573 **self** clause, as described in Section 2.14.4. In C, the arguments are a pointer to the data and  
 2574 length in bytes. In Fortran, two forms are supported. In the first, the argument is a contiguous array  
 2575 section of intrinsic type. In the second, the first argument is a variable or array element and the  
 2576 second is the length in bytes. For data not in shared memory, the data in the local memory is copied  
 2577 to the corresponding device memory. There must be a device copy of the data on the device when  
 2578 calling this routine, otherwise no action is taken by the routine. It is a runtime error to call this  
 2579 routine if the data is not present in the current device memory.

2580 The **\_async** versions of this function will perform the data transfers asynchronously on the **async**  
 2581 queue associated with the value passed in as the **async** argument. The function may return be-  
 2582 fore the data has been transferred; see Section 2.16 Asynchronous Behavior for more details. The  
 2583 synchronous versions will not return until the data has been completely transferred.

### 2584 3.2.26. **acc\_map\_data**

2585 **Summary** The **acc\_map\_data** routine maps previously allocated space in the current device  
 2586 memory to the specified host data.

#### 2587 **Format**

C or C++:

```

void acc_map_data( h_void*, d_void*, size_t );

```

2588 **Description** The **acc\_map\_data** routine is similar to an **enter data** directive with a **create**  
 2589 clause, except instead of allocating new device memory to start a data lifetime, the device address  
 2590 to use for the data lifetime is specified as an argument. The first argument is a host address, fol-  
 2591 lowed by the corresponding device address and the data length in bytes. After this call, when the  
 2592 host data appears in a data clause, the specified device memory will be used. It is an error to call  
 2593 **acc\_map\_data** for host data that is already present in the current device memory. It is undefined  
 2594 to call **acc\_map\_data** with a device address that is already mapped to host data. The device  
 2595 address may be the result of a call to **acc\_malloc**, or may come from some other device-specific  
 2596 API routine. After mapping the device memory, the dynamic reference count for the host data is set  
 2597 to one, but no data movement will occur. Memory mapped by **acc\_map\_data** may not have the  
 2598 associated dynamic reference count decremented to zero, except by a call to **acc\_unmap\_data**.  
 2599 See Section 2.6.6 Reference Counters.

### 2600 **3.2.27. acc\_unmap\_data**

2601 **Summary** The `acc_unmap_data` routine unmaps device data from the specified host data.

#### 2602 **Format**

C or C++:

```
void acc_unmap_data( h_void* );
```

2603 **Description** The `acc_unmap_data` routine is similar to an `exit data` directive with a  
2604 `delete` clause, except the device memory is not deallocated. The argument is pointer to the host  
2605 data. A call to this routine ends the data lifetime for the specified host data. The device memory is  
2606 not deallocated. It is undefined behavior to call `acc_unmap_data` with a host address unless that  
2607 host address was mapped to device memory using `acc_map_data`. After unmapping memory the  
2608 dynamic reference count for the pointer is set to zero, but no data movement will occur. It is an  
2609 error to call `acc_unmap_data` if the structured reference count for the pointer is not zero. See  
2610 Section 2.6.6 Reference Counters.

### 2611 **3.2.28. acc\_deviceptr**

2612 **Summary** The `acc_deviceptr` routine returns the device pointer associated with a specific  
2613 host address.

#### 2614 **Format**

C or C++:

```
d_void* acc_deviceptr( h_void* );
```

2615 **Description** The `acc_deviceptr` routine returns the device pointer associated with a host  
2616 address. The argument is the address of a host variable or array that has an active lifetime on the  
2617 current device. If the data is not present in the current device memory, the routine returns a NULL  
2618 value.

### 2619 **3.2.29. acc\_hostptr**

2620 **Summary** The `acc_hostptr` routine returns the host pointer associated with a specific device  
2621 address.

#### 2622 **Format**

C or C++:

```
h_void* acc_hostptr( d_void* );
```

2623 **Description** The `acc_hostptr` routine returns the host pointer associated with a device ad-  
 2624 dress. The argument is the address of a device variable or array, such as that returned from `acc_deviceptr`,  
 2625 `acc_create` or `acc_copyin`. If the device address is NULL, or does not correspond to any host  
 2626 address, the routine returns a NULL value.

### 2627 3.2.30. `acc_is_present`

2628 **Summary** The `acc_is_present` routine tests whether a variable or array region is accessible  
 2629 from the current device.

#### 2630 **Format**

C or C++:

```
int acc_is_present( h_void*, size_t );
```

Fortran:

```
logical function acc_is_present( a )
logical function acc_is_present( a, len )
type(*), dimension(..) :: a
integer :: len
```

2631 **Description** The `acc_is_present` routine tests whether the specified host data is accessible  
 2632 from the current device. In C, the arguments are a pointer to the data and length in bytes; the  
 2633 function returns nonzero if the specified data is fully present, and zero otherwise. In Fortran, two  
 2634 forms are supported. In the first, the argument is a contiguous array section of intrinsic type. In the  
 2635 second, the first argument is a variable or array element and the second is the length in bytes. The  
 2636 function returns `.true.` if the specified data is in shared memory or is fully present, and `.false.`  
 2637 otherwise. If the byte length is zero, the function returns nonzero in C or `.true.` in Fortran if the  
 2638 given address is in shared memory or is present at all in the current device memory.

### 2639 3.2.31. `acc_memcpy_to_device`

2640 **Summary** The `acc_memcpy_to_device` routine copies data from local memory to device  
 2641 memory.

#### 2642 **Format**

C or C++:

```
void acc_memcpy_to_device( d_void* dest, h_void* src, size_t bytes );
void acc_memcpy_to_device_async( d_void* dest, h_void* src,
size_t bytes, int async );
```

2643 **Description** The `acc_memcpy_to_device` routine copies **bytes** of data from the local  
2644 address in **src** to the device address in **dest**. The destination address must be an address accessible  
2645 from the current device, such as an address returned from `acc_malloc` or `acc_deviceptr`, or  
2646 an address in shared memory.

2647 The `_async` version of this function will perform the data transfers asynchronously on the `async`  
2648 queue associated with the value passed in as the **async** argument. The function may return be-  
2649 fore the data has been transferred; see Section 2.16 Asynchronous Behavior for more details. The  
2650 synchronous versions will not return until the data has been completely transferred.

### 2651 3.2.32. `acc_memcpy_from_device`

2652 **Summary** The `acc_memcpy_from_device` routine copies data from device memory to lo-  
2653 cal memory.

#### 2654 **Format**

C or C++:

```
void acc_memcpy_from_device( h_void* dest, d_void* src, size_t bytes );  
void acc_memcpy_from_device_async( h_void* dest, d_void* src,  
    size_t bytes, int async );
```

2655 **Description** The `acc_memcpy_from_device` routine copies **bytes** data from the device  
2656 address in **src** to the local address in **dest**. The source address must be an address accessible  
2657 from the current device, such as an address returned from `acc_malloc` or `acc_deviceptr`,  
2658 or an address in shared memory.

2659 The `_async` version of this function will perform the data transfers asynchronously on the `async`  
2660 queue associated with the value passed in as the **async** argument. The function may return be-  
2661 fore the data has been transferred; see Section 2.16 Asynchronous Behavior for more details. The  
2662 synchronous versions will not return until the data has been completely transferred.

### 2663 3.2.33. `acc_memcpy_device`

2664 **Summary** The `acc_memcpy_device` routine copies data from one memory location to an-  
2665 other memory location on the current device.

#### 2666 **Format**

C or C++:

```
void acc_memcpy_device( d_void* dest, d_void* src, size_t bytes );  
void acc_memcpy_device_async( d_void* dest, d_void* src,  
    size_t bytes, int async );
```

2667 **Description** The `acc_memcpy_device` routine copies `bytes` data from the device address  
2668 in `src` to the device address in `dest`. Both addresses must be addresses in the current device  
2669 memory, such as would be returned from `acc_malloc` or `acc_deviceptr`. If `dest` and `src`  
2670 overlap, the behavior is undefined.

2671 The `_async` version of this function will perform the data transfers asynchronously on the `async`  
2672 queue associated with the value passed in as the `async` argument. The function may return be-  
2673 fore the data has been transferred; see Section 2.16 Asynchronous Behavior for more details. The  
2674 synchronous versions will not return until the data has been completely transferred.

### 2675 3.2.34. `acc_attach`

2676 **Summary** The `acc_attach` routine updates a pointer in device memory to point to the corre-  
2677 sponding device copy of the host pointer target.

#### 2678 **Format**

C or C++:

```
void acc_attach( h_void** ptr );  
void acc_attach_async( h_void** ptr, int async );
```

2679 **Description** The `acc_attach` routines are passed the address of a host pointer. If the data is  
2680 in shared memory, or if the pointer `*ptr` is in shared memory or is not present in the current device  
2681 memory, or the address to which the `*ptr` points is not present in the current device memory, no  
2682 action is taken. Otherwise, these routines perform the *attach* action (Section 2.7.2).

2683 These routines may issue a data transfer from local memory to device memory. The `_async`  
2684 version of this function will perform the data transfers asynchronously on the `async` queue associated  
2685 with the value passed in as the `async` argument. The function may return before the data has been  
2686 transferred; see Section 2.16 Asynchronous Behavior for more details. The synchronous version  
2687 will not return until the data has been completely transferred.

### 2688 3.2.35. `acc_detach`

2689 **Summary** The `acc_detach` routine updates a pointer in device memory to point to the host  
2690 pointer target.

#### 2691 **Format**

C or C++:

```
void acc_detach( h_void** ptr );  
void acc_detach_async( h_void** ptr, int async );  
void acc_detach_finalize( h_void** ptr );  
void acc_detach_finalize_async( h_void** ptr, int async );
```



2692 **Description** The **acc\_detach** routines are passed the address of a host pointer. If the data is  
2693 in shared memory, or if the pointer **\*ptr** is in shared memory or is not present in the current device  
2694 memory, if the *attachment counter* for the pointer **\*ptr** is zero, no action is taken. Otherwise, these  
2695 routines perform the *detach* action (Section 2.7.2).

2696 The **acc\_detach\_finalize** routines are equivalent to an **exit data** directive with **detach**  
2697 and **finalize** clauses, as described in Section 2.7.12 detach clause. If the data is in shared  
2698 memory, or if the pointer **\*ptr** is not present in the current device memory, or if the *attachment*  
2699 *counter* for the pointer **\*ptr** is zero, no action is taken. Otherwise, these routines perform the  
2700 *immediate detach* action (Section 2.7.2).

2701 These routines may issue a data transfer from local memory to device memory. The **\_async**  
2702 versions of these functions will perform the data transfers asynchronously on the async queue asso-  
2703 ciated with the value passed in as the **async** argument. These functions may return before the data  
2704 has been transferred; see Section 2.16 Asynchronous Behavior for more details. The synchronous  
2705 versions will not return until the data has been completely transferred.



## 2706 4. Environment Variables

2707 This chapter describes the environment variables that modify the behavior of accelerator regions.  
2708 The names of the environment variables must be upper case. The values assigned environment  
2709 variables are case-insensitive and may have leading and trailing white space. If the values of the  
2710 environment variables change after the program has started, even if the program itself modifies the  
2711 values, the behavior is implementation-defined.

### 2712 4.1. ACC\_DEVICE\_TYPE

2713 The **ACC\_DEVICE\_TYPE** environment variable controls the default device type to use when ex-  
2714 ecuting parallel, kernels, and serial regions, if the program has been compiled to use more than  
2715 one different type of device. The allowed values of this environment variable are implementa-  
2716 tion-defined. See the release notes for currently-supported values of this environment variable.

Example:

```
setenv ACC_DEVICE_TYPE NVIDIA
export ACC_DEVICE_TYPE=NVIDIA
```

### 2717 4.2. ACC\_DEVICE\_NUM

2718 The **ACC\_DEVICE\_NUM** environment variable controls the default device number to use when  
2719 executing accelerator regions. The value of this environment variable must be a nonnegative integer  
2720 between zero and the number of devices of the desired type attached to the host. If the value is  
2721 greater than or equal to the number of devices attached, the behavior is implementation-defined.

Example:

```
setenv ACC_DEVICE_NUM 1
export ACC_DEVICE_NUM=1
```

### 2722 4.3. ACC\_PROFLIB

2723 The **ACC\_PROFLIB** environment variable specifies the profiling library. More details about the  
2724 evaluation at runtime is given in section 5.3.3 Runtime Dynamic Library Loading.

Example:

```
setenv ACC_PROFLIB /path/to/proflib/libaccprof.so
export ACC_PROFLIB=/path/to/proflib/libaccprof.so
```



## 5. Profiling Interface

This chapter describes the OpenACC interface for tools that can be used for profile and trace data collection. Therefore it provides a set of OpenACC-specific event callbacks that are triggered during the application run. Currently, this interface does not support tools that employ asynchronous sampling. In this chapter, the term *runtime* refers to the OpenACC runtime library. The term *library* refers to the third party routines invoked at specified events by the OpenACC runtime.

There are four steps for interfacing a *library* to the *runtime*. The first is to write the data collection library callback routines. Section 5.1 Events describes the supported runtime events and the order in which callbacks to the callback routines will occur. Section 5.2 Callbacks Signature describes the signature of the callback routines for all events.

The second is to use registration routines to register the data collection callbacks for the appropriate events. The data collection and registration routines are then saved in a static or dynamic library or shared object. The third is to load the *library* at runtime. The *library* may be statically linked to the application or dynamically loaded by the application or by the *runtime*. This is described in Section 5.3 Loading the Library.

The fourth step is to invoke the registration routine to register the desired callbacks with the events. This may be done explicitly by the application, if the library is statically linked with the application, implicitly by including a call to the registration routine in a `.init` section, or by including an initialization routine in the library if it is dynamically loaded by the *runtime*. This is described in Section 5.4 Registering Event Callbacks.

Subsequently, the *library* may collect information when the callback routines are invoked by the *runtime* and process or store the acquired data.

### 5.1. Events

This section describes the events that are recognized by the runtime. Most events may have a start and end callback routine, that is, a routine that is called just before the runtime code to handle the event starts and another routine that is called just after the event is handled. The event names and routine prototypes are available in the header file `acc_prof.h`, which is delivered with the OpenACC implementation. Event names are prefixed with `acc_ev_`.

The ordering of events must reflect the order in which the OpenACC runtime actually executes them, i.e. if a runtime moves the enqueueing of data transfers or kernel launches outside the originating clauses/constructs, it needs to issue the corresponding launch callbacks when they really occur. A callback for a start event must always precede the matching end callback. The behavior of a tool receiving a callback after the runtime shutdown callback is undefined.

The events that the runtime supports can be registered with a callback and are defined in the enumeration type `acc_event_t`.

```
typedef enum acc_event_t{
    acc_ev_none = 0,
    acc_ev_device_init_start,
    acc_ev_device_init_end,
    acc_ev_device_shutdown_start,
    acc_ev_device_shutdown_end,
    acc_ev_runtime_shutdown,
    acc_ev_create,
    acc_ev_delete,
    acc_ev_alloc,
    acc_ev_free,
    acc_ev_enter_data_start,
    acc_ev_enter_data_end,
    acc_ev_exit_data_start,
    acc_ev_exit_data_end,
    acc_ev_update_start,
    acc_ev_update_end,
    acc_ev_compute_construct_start,
    acc_ev_compute_construct_end,
    acc_ev_enqueue_launch_start,
    acc_ev_enqueue_launch_end,
    acc_ev_enqueue_upload_start,
    acc_ev_enqueue_upload_end,
    acc_ev_enqueue_download_start,
    acc_ev_enqueue_download_end,
    acc_ev_wait_start,
    acc_ev_wait_end,
    acc_ev_last
}acc_event_t;
```

### 2760 5.1.1. Runtime Initialization and Shutdown

2761 No callbacks can be registered for the runtime initialization. Instead the initialization of the tool is  
2762 handled as described in Section 5.3 Loading the Library.

2763 The *runtime shutdown* event name is

```
acc_ev_runtime_shutdown
```

2764 The **acc\_ev\_runtime\_shutdown** event is triggered before the OpenACC runtime shuts down,  
2765 either because all devices have been shutdown by calls to the **acc\_shutdown** API routine, or at  
2766 the end of the program.

### 2767 5.1.2. Device Initialization and Shutdown

2768 The *device initialization* event names are

```
acc_ev_device_init_start
```

**acc\_ev\_device\_init\_end**

2769 These events are triggered when a device is being initialized by the OpenACC runtime. This may be  
2770 when the program starts, or may be later during execution when the program reaches an **acc\_init**  
2771 call or an OpenACC construct. The **acc\_ev\_device\_init\_start** is triggered before device  
2772 initialization starts and **acc\_ev\_device\_init\_end** after initialization is complete.

2773 The *device shutdown* event names are

**acc\_ev\_device\_shutdown\_start****acc\_ev\_device\_shutdown\_end**

2774 These events are triggered when a device is shut down, most likely by a call to the OpenACC  
2775 **acc\_shutdown** API routine. The **acc\_ev\_device\_shutdown\_start** is triggered before  
2776 the device shutdown process starts and **acc\_ev\_device\_shutdown\_end** after the device shut-  
2777 down is complete.

**5.1.3. Enter Data and Exit Data**

2779 The *enter data* and *exit data* event names are

**acc\_ev\_enter\_data\_start****acc\_ev\_enter\_data\_end****acc\_ev\_exit\_data\_start****acc\_ev\_exit\_data\_end**

2780 The **acc\_ev\_enter\_data\_start** and **acc\_ev\_enter\_data\_end** events are triggered at  
2781 **enter data** directives, entry to data constructs, and entry to implicit data regions such as those  
2782 generated by compute constructs. The **acc\_ev\_enter\_data\_start** event is triggered before  
2783 any *data allocation*, *data update*, or *wait* events that are associated with that directive or region  
2784 entry, and the **acc\_ev\_enter\_data\_end** is triggered after those events.

2785 The **acc\_ev\_exit\_data\_start** and **acc\_ev\_exit\_data\_end** events are triggered at **exit**  
2786 **data** directives, exit from **data** constructs, and exit from implicit data regions. The **acc\_ev\_exit\_data\_start**  
2787 event is triggered before any *data deallocation*, *data update*, or *wait* events associated with that di-  
2788 rective or region exit, and the **acc\_ev\_exit\_data\_end** event is triggered after those events.

2789 When the construct that triggers an *enter data* or *exit data* event was generated implicitly by the  
2790 compiler the **implicit** field in the event structure will be set to **1**. When the construct that  
2791 triggers these events was specified explicitly by the application code the **implicit** field in the  
2792 event structure will be set to **0**.

**5.1.4. Data Allocation**

2794 The *data allocation* event names are

**acc\_ev\_create****acc\_ev\_delete**

**acc\_ev\_alloc**  
**acc\_ev\_free**

2795 An **acc\_ev\_alloc** event is triggered when the OpenACC runtime allocates memory from the de-  
2796 vice memory pool, and an **acc\_ev\_free** event is triggered when the runtime frees that memory.  
2797 An **acc\_ev\_create** event is triggered when the OpenACC runtime associates device memory  
2798 with local memory, such as for a data clause (**create**, **copyin**, **copy**, **copyout**) at entry to  
2799 a data construct, compute construct, at an **enter data** directive, or in a call to a data API rou-  
2800 tine (**acc\_copyin**, **acc\_create**, ...). An **acc\_ev\_create** event may be preceded by an  
2801 **acc\_ev\_alloc** event, if newly allocated memory is used for this device data, or it may not, if  
2802 the runtime manages its own memory pool. An **acc\_ev\_delete** event is triggered when the  
2803 OpenACC runtime disassociates device memory from local memory, such as for a data clause at  
2804 exit from a data construct, compute construct, at an **exit data** directive, or in a call to a data API  
2805 routine (**acc\_copyout**, **acc\_delete**, ...). An **acc\_ev\_delete** event may be followed by  
2806 an **acc\_ev\_free** event, if the disassociated device memory is freed, or it may not, if the runtime  
2807 manages its own memory pool.

2808 When the action that generates a *data allocation* event was generated explicitly by the application  
2809 code the **implicit** field in the event structure will be set to **0**. When the *data allocation* event  
2810 is triggered because of a variable or array with implicitly-determined data attributes or otherwise  
2811 implicitly by the compiler the **implicit** field in the event structure will be set to **1**.

### 2812 5.1.5. Data Construct

2813 The events for entering and leaving *data constructs* are mapped to *enter data* and *exit data* events  
2814 as described in Section 5.1.3 Enter Data and Exit Data.

### 2815 5.1.6. Update Directive

2816 The *update directive* event names are

**acc\_ev\_update\_start**  
**acc\_ev\_update\_end**

2817 The **acc\_ev\_update\_start** event will be triggered at an **update** directive, before any *data*  
2818 *update* or *wait* events that are associated with the update directive are carried out, and the corre-  
2819 sponding **acc\_ev\_update\_end** event will be triggered after any of the associated events.

### 2820 5.1.7. Compute Construct

2821 The *compute construct* event names are

**acc\_ev\_compute\_construct\_start**  
**acc\_ev\_compute\_construct\_end**

2822 The **acc\_ev\_compute\_construct\_start** event is triggered at entry to a compute construct,  
2823 before any *launch* events that are associated with entry to the compute construct. The **acc\_ev\_compute\_construct**



2824 event is triggered at the exit of the compute construct, after any *launch* events associated with exit  
2825 from the compute construct. If there are data clauses on the compute construct, those data clauses  
2826 may be treated as part of the compute construct, or as part of a data construct containing the compute  
2827 construct. The callbacks for data clauses must use the same line numbers as for the compute  
2828 construct events.

### 2829 5.1.8. Enqueue Kernel Launch

2830 The *launch* event names are

```
acc_ev_enqueue_launch_start
acc_ev_enqueue_launch_end
```

2831 The **acc\_ev\_enqueue\_launch\_start** event is triggered just before an accelerator compu-  
2832 tation is enqueued for execution on a device, and **acc\_ev\_enqueue\_launch\_end** is trig-  
2833 gered just after the computation is enqueued. Note that these events are synchronous with the  
2834 local thread enqueueing the computation to a device, not with the device executing the compu-  
2835 tation. The **acc\_ev\_enqueue\_launch\_start** event callback routine is invoked just before  
2836 the computation is enqueued, not just before the computation starts execution. More importantly,  
2837 the **acc\_ev\_enqueue\_launch\_end** event callback routine is invoked after the computation is  
2838 enqueued, not after the computation finished executing.

2839 **Note:** Measuring the time between the start and end launch callbacks is often unlikely to be useful,  
2840 since it will only measure the time to manage the launch queue, not the time to execute the code on  
2841 the device.

### 2842 5.1.9. Enqueue Data Update (Upload and Download)

2843 The *data update* event names are

```
acc_ev_enqueue_upload_start
acc_ev_enqueue_upload_end
acc_ev_enqueue_download_start
acc_ev_enqueue_download_end
```

2844 The **\_start** events are triggered just before each upload (data copy from local memory to device  
2845 memory) operation is or download (data copy from device memory to local memory) operation is  
2846 enqueued for execution on a device. The corresponding **\_end** events are triggered just after each  
2847 upload or download operation is enqueued.

2848 **Note:** Measuring the time between the start and end update callbacks is often unlikely to be useful,  
2849 since it will only measure the time to manage the enqueue operation, not the time to perform the  
2850 actual upload or download.

2851 When the action that generates a *data update* event was generated explicitly by the application  
2852 code the **implicit** field in the event structure will be set to **0**. When the *data allocation* event  
2853 is triggered because of a variable or array with implicitly-determined data attributes or otherwise  
2854 implicitly by the compiler the **implicit** field in the event structure will be set to **1**.

2855 **5.1.10. Wait**2856 The *wait* event names are

```

acc_ev_wait_start
acc_ev_wait_end

```

2857 An **acc\_ev\_wait\_start** will be triggered for each relevant queue before the local thread waits  
 2858 for that queue to be empty. A **acc\_ev\_wait\_end** will be triggered for each relevant queue after  
 2859 the local thread has determined that the queue is empty.

2860 Wait events occur when the local thread and a device synchronize, either due to a **wait** directive  
 2861 or by a *wait* clause on a synchronous data construct, compute construct, or **enter data**, **exit**  
 2862 **data**, or **update** directive. For *wait* events triggered by an explicit synchronous **wait** directive  
 2863 or *wait* clause, the **implicit** field in the event structure will be **0**. For all other wait events, the  
 2864 **implicit** field in the event structure will be **1**.

2865 The OpenACC runtime need not trigger *wait* events for queues that have not been used in the  
 2866 program, and need not trigger *wait* events for queues that have not been used by this thread since  
 2867 the last *wait* operation. For instance, an **acc wait** directive with no arguments is defined to wait on  
 2868 all queues. If the program only uses the default (synchronous) queue and the queue associated with  
 2869 **async(1)** and **async(2)** then an **acc wait** directive may trigger *wait* events only for those  
 2870 three queues. If the implementation knows that no activities have been enqueued on the **async(2)**  
 2871 queue since the last *wait* operation, then the **acc wait** directive may trigger *wait* events only for  
 2872 the default queue and the **async(1)** queue.

2873 **5.2. Callbacks Signature**

2874 This section describes the signature of event callbacks. All event callbacks have the same signature.  
 2875 The routine prototypes are available in the header file **acc\_prof.h**, which is delivered with the  
 2876 OpenACC implementation.

2877 All callback routines have three arguments. The first argument is a pointer to a struct containing  
 2878 general information; the same struct type is used for all callback events. The second argument is  
 2879 a pointer to a struct containing information specific to that callback event; there is one struct type  
 2880 containing information for data events, another struct type containing information for kernel launch  
 2881 events, and a third struct type for other events, containing essentially no information. The third  
 2882 argument is a pointer to a struct containing information about the application programming interface  
 2883 (API) being used for the specific device. For NVIDIA CUDA devices, this contains CUDA-specific  
 2884 information; for OpenCL devices, this contains OpenCL-specific information. Other interfaces can  
 2885 be supported as they are added by implementations. The prototype for a callback routine is:

```

typedef void (*acc_prof_callback)
(acc_prof_info*, acc_event_info*, acc_api_info*);

```

2886 In the descriptions, the datatype **ssize\_t** means a signed 32-bit integer for a 32-bit binary and  
 2887 a 64-bit integer for a 64-bit binary, the datatype **size\_t** means an unsigned 32-bit integer for a

2888 32-bit binary and a 64-bit integer for a 64-bit binary, and the datatype **int** means a 32-bit integer  
 2889 for both 32-bit and 64-bit binaries. A null pointer is the pointer with value zero.

### 2890 5.2.1. First Argument: General Information

2891 The first argument is a pointer to the **acc\_prof\_info** struct type:

```
typedef struct acc_prof_info{
    acc_event_t event_type;
    int valid_bytes;
    int version;
    acc_device_t device_type;
    int device_number;
    int thread_id;
    ssize_t async;
    ssize_t async_queue;
    const char* src_file;
    const char* func_name;
    int line_no, end_line_no;
    int func_line_no, func_end_line_no;
}acc_prof_info;
```

2892 The fields are described below.

2893 • **acc\_event\_t event\_type** - The event type that triggered this callback. The datatype  
 2894 is the enumeration type **acc\_event\_t**, described in the previous section. This allows the  
 2895 same callback routine to be used for different events.

2896 • **int valid\_bytes** - The number of valid bytes in this struct. This allows a library to inter-  
 2897 face with newer runtimes that may add new fields to the struct at the end while retaining com-  
 2898 patibility with older runtimes. A runtime must fill in the **event\_type** and **valid\_bytes**  
 2899 fields, and must fill in values for all fields with offset less than **valid\_bytes**. The value of  
 2900 **valid\_bytes** for a struct is recursively defined as:

```
valid_bytes(struct) = offset(lastfield) + valid_bytes(lastfield)
valid_bytes(type[n]) = (n-1)*sizeof(type) + valid_bytes(type)
valid_bytes(basictype) = sizeof(basictype)
```

2901 • **int version** - A version number; the value of **\_OPENACC**.

2902 • **acc\_device\_t device\_type** - The device type corresponding to this event. The datatype  
 2903 is **acc\_device\_t**, an enumeration type of all the supported device types, defined in **openacc.h**.

2904 • **int device\_number** - The device number. Each device is numbered, typically starting at  
 2905 device zero. For applications that use more than one device type, the device numbers may be  
 2906 unique across all devices or may be unique only across all devices of the same device type.

2907 • **int thread\_id** - The host thread ID making the callback. Host threads are given unique  
 2908 thread ID numbers typically starting at zero. This is not necessarily the same as the OpenMP  
 2909 thread number.

- 2910 • **ssize\_t async** - The value of the **async()** clause for the directive that triggered this  
2911 callback.
- 2912 • **ssize\_t async\_queue** - If the runtime uses a limited number of asynchronous queues,  
2913 this field contains the internal asynchronous queue number used for the event.
- 2914 • **const char\* src\_file** - A pointer to null-terminated string containing the name of or  
2915 path to the source file, if known, or a null pointer if not. If the library wants to save the source  
2916 file name, it should allocate memory and copy the string.
- 2917 • **const char\* func\_name** - A pointer to a null-terminated string containing the name of  
2918 the function in which the event occurred, if known, or a null pointer if not. If the library wants  
2919 to save the function name, it should allocate memory and copy the string.
- 2920 • **int line\_no** - The line number of the directive or program construct or the starting line  
2921 number of the OpenACC construct corresponding to the event. A negative or zero value  
2922 means the line number is not known.
- 2923 • **int end\_line\_no** - For an OpenACC construct, this contains the line number of the end  
2924 of the construct. A negative or zero value means the line number is not known.
- 2925 • **int func\_line\_no** - The line number of the first line of the function named in **func\_name**.  
2926 A negative or zero value means the line number is not known.
- 2927 • **int func\_end\_line\_no** - The last line number of the function named in **func\_name**.  
2928 A negative or zero value means the line number is not known.

### 2929 5.2.2. Second Argument: Event-Specific Information

2930 The second argument is a pointer to the **acc\_event\_info** union type.

```

2931 typedef union acc_event_info{
2932     acc_event_t event_type;
2933     acc_data_event_info data_event;
2934     acc_launch_event_info launch_event;
2935     acc_other_event_info other_event;
2936 }acc_event_info;

```

2931 The **event\_type** field selects which union member to use. The first five members of each union  
2932 member are identical. The second through fifth members of each union member (**valid\_bytes**,  
2933 **parent\_construct**, **implicit**, and **tool\_info**) have the same semantics for all event  
2934 types:

- 2935 • **int valid\_bytes** - The number of valid bytes in the respective struct. (This field is similar  
2936 used as discussed in Section 5.2.1 First Argument: General Information.)
- 2937 • **acc\_construct\_t parent\_construct** - This field describes the type of construct  
2938 that caused the event to be emitted. The possible values for this field are defined by the  
2939 **acc\_construct\_t** enum, described at the end of this section.
- 2940 • **int implicit** - This field is set to 1 for any implicit event, such as an implicit wait at  
2941 a synchronous data construct or synchronous enter data, exit data or update directive. This

2942 field is set to zero when the event is triggered by an explicit directive or call to a runtime API  
 2943 routine.

2944 • **void\* tool\_info** - This field is used to pass tool-specific information from a **\_start**  
 2945 event to the matching **\_end** event. For a **\_start** event callback, this field will be initialized  
 2946 to a null pointer. The value of this field for a **\_end** event will be the value returned by  
 2947 the library in this field from the matching **\_start** event callback, if there was one, or null  
 2948 otherwise. For events that are neither **\_start** or **\_end** events, this field will be null.

## 2949 Data Events

2950 For a data event, as noted in the event descriptions, the second argument will be a pointer to the  
 2951 **acc\_data\_event\_info** struct.

```

typedef struct acc_data_event_info{
    acc_event_t event_type;
    int valid_bytes;
    acc_construct_t parent_construct;
    int implicit;
    void* tool_info;
    const char* var_name;
    size_t bytes;
    const void* host_ptr;
    const void* device_ptr;
}acc_data_event_info;
  
```

2952 The fields specific for a data event are:

2953 • **acc\_event\_t event\_type** - The event type that triggered this callback. The events that  
 2954 use the **acc\_data\_event\_info** struct are:

```

    acc_ev_enqueue_upload_start
    acc_ev_enqueue_upload_end
    acc_ev_enqueue_download_start
    acc_ev_enqueue_download_end
    acc_ev_create
    acc_ev_delete
    acc_ev_alloc
    acc_ev_free
  
```

2955 • **const char\* var\_name** - A pointer to null-terminated string containing the name of the  
 2956 variable for which this event is triggered, if known, or a null pointer if not. If the library wants  
 2957 to save the variable name, it should allocate memory and copy the string.

2958 • **size\_t bytes** - The number of bytes for the data event.

2959 • **const void\* host\_ptr** - If available and appropriate for this event, this is a pointer to  
 2960 the host data.

2961 • **const void\* device\_ptr** - If available and appropriate for this event, this is a pointer  
 2962 to the corresponding device data.

2963 **Launch Events**

2964 For a launch event, as noted in the event descriptions, the second argument will be a pointer to the  
 2965 **acc\_launch\_event\_info** struct.

```

typedef struct acc_launch_event_info{
    acc_event_t event_type;
    int valid_bytes;
    acc_construct_t parent_construct;
    int implicit;
    void* tool_info;
    const char* kernel_name;
    size_t num_gangs, num_workers, vector_length;
}acc_launch_event_info;

```

2966 The fields specific for a launch event are:

2967 • **acc\_event\_t event\_type** - The event type that triggered this callback. The events that  
 2968 use the **acc\_launch\_event\_info** struct are:

```

    acc_ev_enqueue_launch_start
    acc_ev_enqueue_launch_end

```

2969 • **const char\* kernel\_name** - A pointer to null-terminated string containing the name of  
 2970 the kernel being launched, if known, or a null pointer if not. If the library wants to save the  
 2971 kernel name, it should allocate memory and copy the string.

2972 • **size\_t num\_gangs, num\_workers, vector\_length** - The number of gangs, work-  
 2973 ers and vector lanes created for this kernel launch.

2974 **Other Events**

2975 For any event that does not use the **acc\_data\_event\_info** or **acc\_launch\_event\_info**  
 2976 struct, the second argument to the callback routine will be a pointer to **acc\_other\_event\_info**  
 2977 struct.

```

typedef struct acc_other_event_info{
    acc_event_t event_type;
    int valid_bytes;
    acc_construct_t parent_construct;
    int implicit;
    void* tool_info;
}acc_other_event_info;

```

2978 **Parent Construct Enumeration**

2979 All event structures contain a **parent\_construct** member that describes the type of construct  
 2980 that caused the event to be emitted. The purpose of this field is to provide a means to identify

2981 the type of construct emitting the event in the cases where an event may be emitted by multi-  
 2982 ple construct types, such as is the case with data and wait events. The possible values for the  
 2983 **parent\_construct** field are defined in the enumeration type **acc\_construct\_t**. In the  
 2984 case of combined directives, the outermost construct of the combined construct should be specified  
 2985 as the **parent\_construct**. If the event was emitted as the result of the application making a  
 2986 call to the runtime api, the value will be **acc\_construct\_runtime\_api**.

```

typedef enum acc_construct_t{
    acc_construct_parallel = 0,
    acc_construct_kernels,
    acc_construct_loop,
    acc_construct_data,
    acc_construct_enter_data,
    acc_construct_exit_data,
    acc_construct_host_data,
    acc_construct_atomic,
    acc_construct_declare,
    acc_construct_init,
    acc_construct_shutdown,
    acc_construct_set,
    acc_construct_update,
    acc_construct_routine,
    acc_construct_wait,
    acc_construct_runtime_api,
    acc_construct_serial
}acc_construct_t;

```

### 2987 5.2.3. Third Argument: API-Specific Information

2988 The third argument is a pointer to the **acc\_api\_info** struct type, shown here.

```

typedef union acc_api_info{
    acc_device_api device_api;
    int valid_bytes;
    acc_device_t device_type;
    int vendor;
    const void* device_handle;
    const void* context_handle;
    const void* async_handle;
}acc_api_info;

```

2989 The fields are described below:

- 2990 • **acc\_device\_api device\_api** - The API in use for this device. The data type is the  
 2991 enumeration **acc\_device\_api**, which is described later in this section.
- 2992 • **int valid\_bytes** - The number of valid bytes in this struct. See the discussion above in  
 2993 Section 5.2.1 First Argument: General Information.

- 2994 • **acc\_device\_t device\_type** - The device type; the datatype is **acc\_device\_t**, de-  
2995 fined in **openacc.h**.
- 2996 • **int vendor** - An identifier to identify the OpenACC vendor; contact your vendor to deter-  
2997 mine the value used by that vendor's runtime.
- 2998 • **const void\* device\_handle** - If applicable, this will be a pointer to the API-specific  
2999 device information.
- 3000 • **const void\* context\_handle** - If applicable, this will be a pointer to the API-specific  
3001 context information.
- 3002 • **const void\* async\_handle** - If applicable, this will be a pointer to the API-specific  
3003 async queue information.

3004 According to the value of **device\_api** a library can cast the pointers of the fields **device\_handle**,  
3005 **context\_handle** and **async\_handle** to the respective device API type. The following device  
3006 APIs are defined in this interface:

```

typedef enum acc_device_api{
    acc_device_api_none = 0,    /* no device API */
    acc_device_api_cuda,       /* CUDA driver API */
    acc_device_api_opencl,     /* OpenCL API */
    acc_device_api_coi,        /* COI API */
    acc_device_api_other       /* other device API */
}acc_device_api;

```

### 3007 5.3. Loading the Library

3008 This section describes how a tools library is loaded when the program is run. Four methods are  
3009 described.

- 3010 • A tools library may be linked with the program, as any other library is linked, either as a  
3011 static library or a dynamic library, and the runtime will call a predefined library initialization  
3012 routine that will register the event callbacks.
- 3013 • The OpenACC runtime implementation may support a dynamic tools library, such as a shared  
3014 object for Linux or OS/X, or a DLL for Windows, which is then dynamically loaded at runtime  
3015 under control of the environment variable **ACC\_PROFLIB**.
- 3016 • Some implementations where the OpenACC runtime is itself implemented as a dynamic li-  
3017 brary may support adding a tools library using the **LD\_PRELOAD** feature in Linux.
- 3018 • A tools library may be linked with the program, as in the first option, and the application itself  
3019 can call a library initialization routine that will register the event callbacks.

3020 Callbacks are registered with the runtime by calling **acc\_prof\_register** for each event as  
3021 described in Section 5.4 Registering Event Callbacks. The prototype for **acc\_prof\_register**  
3022 is:

```

extern void acc_prof_register
    (acc_event_t event_type, acc_prof_callback cb,

```



```
acc_register_t info);
```

3023 The first argument to **acc\_prof\_register** is the event for which a callback is being registered  
3024 (compare Section 5.1 Events). The second argument is a pointer to the callback routine:

```
typedef void (*acc_prof_callback)
    (acc_prof_info*, acc_event_info*, acc_api_info*);
```

3025 The third argument is usually zero (or **acc\_reg**). See Section 5.4.2 Disabling and Enabling Callbacks  
3026 for cases where a nonzero value is used. The argument **acc\_register\_t** is an enum type:

```
typedef enum acc_register_t{
    acc_reg = 0,
    acc_toggle = 1,
    acc_toggle_per_thread = 2
}acc_register_t;
```

3027 An example of registering callbacks for launch, upload, and download events is:

```
acc_prof_register(acc_ev_enqueue_launch_start, prof_launch, 0);
acc_prof_register(acc_ev_enqueue_upload_start, prof_data, 0);
acc_prof_register(acc_ev_enqueue_download_start, prof_data, 0);
```

3028 As shown in this example, the same routine (**prof\_data**) can be registered for multiple events.  
3029 The routine can use the **event\_type** field in the **acc\_prof\_info** structure to determine for  
3030 what event it was invoked.

### 3031 5.3.1. Library Registration

3032 The OpenACC runtime will invoke **acc\_register\_library**, passing the addresses of the reg-  
3033 istration routines **acc\_prof\_register** and **acc\_prof\_unregister**, in case that routine  
3034 comes from a dynamic library. In the third argument it passes the address of the lookup routine  
3035 **acc\_prof\_lookup** to obtain the addresses of inquiry functions. No inquiry functions are de-  
3036 fined in this profiling interface, but we preserve this argument for future support of sampling-based  
3037 tools.

3038 Typically, the OpenACC runtime will include a *weak* definition of **acc\_register\_library**,  
3039 which does nothing and which will be called when there is no tools library. In this case, the library  
3040 can save the addresses of these routines and/or make registration calls to register any appropriate  
3041 callbacks. The prototype for **acc\_register\_library** is:

```
extern void acc_register_library
    (acc_prof_reg register, acc_prof_reg unregister,
    acc_prof_lookup_func lookup);
```

3042 The first two arguments of this routine are of type:

```

typedef void (*acc_prof_reg)
    (acc_event_t event_type, acc_prof_callback cb,
     acc_register_t info);

```

3043 The third argument passes the address to the lookup function `acc_prof_lookup` to obtain the  
 3044 address of interface functions. It is of type:

```

typedef void (*acc_query_fn) ();
typedef acc_query_fn (*acc_prof_lookup_func)
    (const char* acc_query_fn_name);

```

3045 The argument of the lookup function is a string with the name of the inquiry function. There are no  
 3046 inquiry functions defined for this interface.

### 3047 5.3.2. Statically-Linked Library Initialization

3048 A tools library can be compiled and linked directly into the application. If the library provides an  
 3049 external routine `acc_register_library` as specified in Section 5.3.1 Library Registration, the  
 3050 runtime will invoke that routine to initialize the library.

3051 The sequence of events is:

- 3052 1. The runtime invokes the `acc_register_library` routine from the library.
- 3053 2. The `acc_register_library` routine calls `acc_prof_register` for each event to  
 3054 be monitored.
- 3055 3. `acc_prof_register` records the callback routines.
- 3056 4. The program runs, and your callback routines are invoked at the appropriate events.

3057 In this mode, only one tool library is supported.

### 3058 5.3.3. Runtime Dynamic Library Loading

3059 A common case is to build the tools library as a dynamic library (shared object for Linux or OS/X,  
 3060 DLL for Windows). In that case, you can have the OpenACC runtime load the library during initial-  
 3061 ization. This allows you to enable runtime profiling without rebuilding or even relinking your ap-  
 3062 plication. The dynamic library must implement a registration routine `acc_register_library`  
 3063 as specified in Section 5.3.1 Library Registration.

3064 The user may set the environment variable `ACC_PROFLIB` to the path to the library will tell the  
 3065 OpenACC runtime to load your dynamic library at initialization time:

```

Bash:
    export ACC_PROFLIB=/home/user/lib/myprof.so
    ./myapp
or
    ACC_PROFLIB=/home/user/lib/myprof.so ./myapp

```

C-shell:

```
setenv ACC_PROFLIB /home/user/lib/myprof.so
./myapp
```

3066 When the OpenACC runtime initializes, it will read the **ACC\_PROFLIB** environment variable (with  
3067 **getenv**). The runtime will open the dynamic library (using **dlopen** or **LoadLibraryA**); if  
3068 the library cannot be opened, the runtime may abort, or may continue execution with or with-  
3069 out an error message. If the library is successfully opened, the runtime will get the address of  
3070 the **acc\_register\_library** routine (using **dlsym** or **GetProcAddress**). If this routine  
3071 is resolved in the library, it will be invoked passing in the addresses of the registration routine  
3072 **acc\_prof\_register**, the deregistration routine **acc\_prof\_unregister**, and the lookup  
3073 routine **acc\_prof\_lookup**. The registration routine in your library, **acc\_register\_library**,  
3074 should register the callbacks by calling the **register** argument, and should save the addresses of  
3075 the arguments (**register**, **unregister**, and **lookup**) for later use, if needed.

3076 The sequence of events is:

- 3077 1. Initialization of the OpenACC runtime.
- 3078 2. OpenACC runtime reads **ACC\_PROFLIB**.
- 3079 3. OpenACC runtime loads the library.
- 3080 4. OpenACC runtime calls the **acc\_register\_library** routine in that library.
- 3081 5. Your **acc\_register\_library** routine calls **acc\_prof\_register** for each event to  
3082 be monitored.
- 3083 6. **acc\_prof\_register** records the callback routines.
- 3084 7. The program runs, and your callback routines are invoked at the appropriate events.

3085 If supported, paths to multiple dynamic libraries may be specified in the **ACC\_PROFLIB** environ-  
3086 ment variable, separated by semicolons (;). The OpenACC runtime will open these libraries and in-  
3087 voke the **acc\_register\_library** routine for each, in the order they appear in **ACC\_PROFLIB**.

#### 3088 5.3.4. Preloading with LD\_PRELOAD

3089 The implementation may also support dynamic loading of a tools library using the **LD\_PRELOAD**  
3090 feature available in some systems. In such an implementation, you need only specify your tools  
3091 library path in the **LD\_PRELOAD** environment variable before executing your program. The Open-  
3092 ACC runtime will invoke the **acc\_register\_library** routine in your tools library at initial-  
3093 ization time. This requires that the OpenACC runtime include a dynamic library with a default  
3094 (empty) implementation of **acc\_register\_library** that will be invoked in the normal case  
3095 where there is no **LD\_PRELOAD** setting. If an implementation only supports static linking, or if the  
3096 application is linked without dynamic library support, this feature will not be available.

Bash:

```
export LD_PRELOAD=/home/user/lib/myprof.so
./myapp
```

or

```
LD_PRELOAD=/home/user/lib/myprof.so ./myapp
```

C-shell:

```
setenv LD_PRELOAD /home/user/lib/myprof.so
./myapp
```

3097 The sequence of events is:

- 3098 1. The operating system loader loads the library specified in **LD\_PRELOAD**.
- 3099 2. The call to **acc\_register\_library** in the OpenACC runtime is resolved to the routine  
3100 in the loaded tools library.
- 3101 3. OpenACC runtime calls the **acc\_register\_library** routine in that library.
- 3102 4. Your **acc\_register\_library** routine calls **acc\_prof\_register** for each event to  
3103 be monitored.
- 3104 5. **acc\_prof\_register** records the callback routines.
- 3105 6. The program runs, and your callback routines are invoked at the appropriate events.

3106 In this mode, only a single tools library is supported, since only one **acc\_register\_library**  
3107 initialization routine will get resolved by the dynamic loader.

### 3108 5.3.5. Application-Controlled Initialization

3109 An alternative to default initialization is to have the application itself call the library initialization  
3110 routine, which then calls **acc\_prof\_register** for each appropriate event. The library may be  
3111 statically linked to the application or your application may dynamically load the library.

3112 The sequence of events is:

- 3113 1. Your application calls the library initialization routine.
- 3114 2. The library initialization routine calls **acc\_prof\_register** for each event to be moni-  
3115 tored.
- 3116 3. **acc\_prof\_register** records the callback routines.
- 3117 4. The program runs, and your callback routines are invoked at the appropriate events.

3118 In this mode, multiple tools libraries can be supported, with each library initialization routine in-  
3119 voked by the application.

## 3120 5.4. Registering Event Callbacks

3121 This section describes how to register and unregister callbacks, temporarily disabling and enabling  
3122 callbacks, the behavior of dynamic registration and unregistration, and requirements on an Open-  
3123 ACC implementation to correctly support the interface.

3124 **5.4.1. Event Registration and Unregistration**

3125 The library must call the registration routine `acc_prof_register` to register each callback  
 3126 with the runtime. A simple example:

```
extern void prof_data(acc_prof_info* profinfo,
                    acc_event_info* eventinfo, acc_api_info* apiinfo);
extern void prof_launch(acc_prof_info* profinfo,
                      acc_event_info* eventinfo, acc_api_info* apiinfo);
...
void acc_register_library(){
    acc_prof_register(acc_ev_enqueue_upload_start, prof_data, 0);
    acc_prof_register(acc_ev_enqueue_download_start, prof_data, 0);
    acc_prof_register(acc_ev_enqueue_launch_start, prof_launch, 0);
}
```

3127 In this example the `prof_data` routine will be invoked for each data upload and download event,  
 3128 and the `prof_launch` routine will be invoked for each launch event. The `prof_data` routine  
 3129 might start out with:

```
void prof_data(acc_prof_info* profinfo,
              acc_event_info* eventinfo, acc_api_info* apiinfo){
    acc_data_event_info* datainfo;
    datainfo = (acc_data_event_info*)eventinfo;
    switch( datainfo->event_type ){
        case acc_ev_enqueue_upload_start :
            ...
    }
}
```

3130 **Multiple Callbacks**

3131 Multiple callback routines can be registered on the same event:

```
acc_prof_register(acc_ev_enqueue_upload_start, prof_data, 0);
acc_prof_register(acc_ev_enqueue_upload_start, prof_up, 0);
```

3132 For most events, the callbacks will be invoked in the order in which they are registered. However,  
 3133 *end* events, named `acc_ev_..._end`, invoke callbacks in the reverse order. Essentially, each  
 3134 event has an ordered list of callback routines. A new callback routine is appended to the tail of the  
 3135 list for that event. For most events, that list is traversed from the head to the tail, but for *end* events,  
 3136 the list is traversed from the tail to the head.

3137 If a callback is registered, then later unregistered, then later still registered again, the second regis-  
 3138 tration is considered to be a new callback, and the callback routine will then be appended to the tail  
 3139 of the callback list for that event.

3140 **Unregistering**

3141 A matching call to **acc\_prof\_unregister** will remove that routine from the list of callback  
 3142 routines for that event.

```

acc_prof_register(acc_ev_enqueue_upload_start, prof_data, 0);
// prof_data is on the callback list for acc_ev_enqueue_upload_start
...
acc_prof_unregister(acc_ev_enqueue_upload_start, prof_data, 0);
// prof_data is removed from the callback list
// for acc_ev_enqueue_upload_start

```

3143 Each entry on the callback list must also have a *ref* count. This keeps track of how many times  
 3144 this routine was added to this event's callback list. If a routine is registered *n* times, it must be  
 3145 unregistered *n* times before it is removed from the list. Note that if a routine is registered multiple  
 3146 times for the same event, its *ref* count will be incremented with each registration, but it will only be  
 3147 invoked once for each event instance.

3148 **5.4.2. Disabling and Enabling Callbacks**

3149 A callback routine may be temporarily disabled on the callback list for an event, then later re-  
 3150 enabled. The behavior is slightly different than unregistering and later re-registering that event.  
 3151 When a routine is disabled and later re-enabled, the routine's position on the callback list for that  
 3152 event is preserved. When a routine is unregistered and later re-registered, the routine's position on  
 3153 the callback list for that event will move to the tail of the list. Also, unregistering a callback must be  
 3154 done *n* times if the callback routine was registered *n* times. In contrast, disabling, and enabling an  
 3155 event sets a toggle. Disabling a callback will immediately reset the toggle and disable calls to that  
 3156 routine for that event, even if it was enabled multiple times. Enabling a callback will immediately  
 3157 set the toggle and enable calls to that routine for that event, even if it was disabled multiple times.  
 3158 Registering a new callback initially sets the toggle.

3159 A call to **acc\_prof\_unregister** with a value of **acc\_toggle** as the third argument will dis-  
 3160 able callbacks to the given routine. A call to **acc\_prof\_register** with a value of **acc\_toggle**  
 3161 as the third argument will enable those callbacks.

```

acc_prof_unregister(acc_ev_enqueue_upload_start,
    prof_data, acc_toggle);
// prof_data is disabled
...
acc_prof_register(acc_ev_enqueue_upload_start,
    prof_data, acc_toggle);
// prof_data is re-enabled

```

3162 A call to either **acc\_prof\_unregister** or **acc\_prof\_register** to disable or enable a call-  
 3163 back when that callback is not currently registered for that event will be ignored with no error.

3164 All callbacks for an event may be disabled (and re-enabled) by passing **NULL** to the second argument  
 3165 and **acc\_toggle** to the third argument of **acc\_prof\_unregister** (and **acc\_prof\_register**).

3166 This sets a toggle for that event, which is distinct from the toggle for each callback for that event.  
 3167 While the event is disabled, no callbacks for that event will be invoked. Callbacks for that event can  
 3168 be registered, unregistered, enabled, and disabled while that event is disabled, but no callbacks will  
 3169 be invoked for that event until the event itself is enabled. Initially, all events are enabled.

```

acc_prof_unregister(acc_ev_enqueue_upload_start,
    prof_data, acc_toggle);
// prof_data is disabled
...
acc_prof_unregister(acc_ev_enqueue_upload_start,
    NULL, acc_toggle);
// acc_ev_enqueue_upload_start callbacks are disabled
...
acc_prof_register(acc_ev_enqueue_upload_start,
    prof_data, acc_toggle);
// prof_data is re-enabled, but
// acc_ev_enqueue_upload_start callbacks still disabled
...
acc_prof_register(acc_ev_enqueue_upload_start, prof_up, 0);
// prof_up is registered and initially enabled, but
// acc_ev_enqueue_upload_start callbacks still disabled
...
acc_prof_register(acc_ev_enqueue_upload_start,
    NULL, acc_toggle);
// acc_ev_enqueue_upload_start callbacks are enabled

```

3170 Finally, all callbacks can be disabled (and enabled) by passing the argument list **(0, NULL,**  
 3171 **acc\_toggle)** to **acc\_prof\_unregister** (and **acc\_prof\_register**). This sets a global  
 3172 toggle disabling all callbacks, which is distinct from the toggle enabling callbacks for each event and  
 3173 the toggle enabling each callback routine. The behavior of passing zero as the first argument and a  
 3174 non-**NULL** value as the second argument to **acc\_prof\_unregister** or **acc\_prof\_register**  
 3175 is not defined, and may be ignored by the runtime without error.

3176 All callbacks can be disabled (or enabled) for just the current thread by passing the argument list  
 3177 **(0, NULL, acc\_toggle\_per\_thread)** to **acc\_prof\_unregister** (and **acc\_prof\_register**).  
 3178 This is the only thread-specific interface to **acc\_prof\_register** and **acc\_prof\_unregister**,  
 3179 all other calls to register, unregister, enable, or disable callbacks affect all threads in the application.

## 3180 5.5. Advanced Topics

3181 This section describes advanced topics such as dynamic registration and changes of the execution  
 3182 state for callback routines as well as the runtime and tool behavior for multiple host threads.

### 3183 5.5.1. Dynamic Behavior

3184 Callback routines may be registered or unregistered, enabled or disabled at any point in the execution  
3185 of the program. Calls may appear in the library itself, during the processing of an event. The  
3186 OpenACC runtime must allow for this case, where the callback list for an event is modified while  
3187 that event is being processed.

#### 3188 Dynamic Registration and Unregistration

3189 Calls to **acc\_register** and **acc\_unregister** may occur at any point in the application. A  
3190 callback routine can be registered or unregistered from a callback routine, either the same routine  
3191 or another routine, for a different event or the same event for which the callback was invoked. If a  
3192 callback routine is registered for an event while that event is being processed, then the new callback  
3193 routine will be added to the tail of the list of callback routines for this event. Some events (the  
3194 **\_end**) events process the callback routines in reverse order, from the tail to the head. For those  
3195 events, adding a new callback routine will not cause the new routine to be invoked for this instance  
3196 of the event. The other events process the callback routines in registration order, from the head to  
3197 the tail. Adding a new callback routine for such a event will cause the runtime to invoke that newly  
3198 registered callback routine for this instance of the event. Both the runtime and the library must  
3199 implement and expect this behavior.

3200 If an existing callback routine is unregistered for an event while that event is being processed, that  
3201 callback routine is removed from the list of callbacks for this event. For any event, if that callback  
3202 routine had not yet been invoked for this instance of the event, it will not be invoked.

3203 Registering and unregistering a callback routine is a global operation and affects all threads, in a  
3204 multithreaded application. See Section 5.4.1 Multiple Callbacks.

#### 3205 Dynamic Enabling and Disabling

3206 Calls to **acc\_register** and **acc\_unregister** to enable and disable a specific callback for  
3207 an event, enable or disable all callbacks for an event, or enable or disable all callbacks may occur  
3208 at any point in the application. A callback routine can be enabled or disabled from a callback  
3209 routine, either the same routine or another routine, for a different event or the same event for which  
3210 the callback was invoked. If a callback routine is enabled for an event while that event is being  
3211 processed, then the new callback routine will be immediately enabled. If it appears on the list of  
3212 callback routines closer to the head (for **\_end** events) or closer to the tail (for other events), that  
3213 newly-enabled callback routine will be invoked for this instance of this event, unless it is disabled  
3214 or unregistered before that callback is reached.

3215 If a callback routine is disabled for an event while that event is being processed, that callback routine  
3216 is immediately disabled. For any event, if that callback routine had not yet been invoked for this in-  
3217 stance of the event, it will not be invoked, unless it is enabled before that callback routine is reached  
3218 in the list of callbacks for this event. If all callbacks for an event are disabled while that event is  
3219 being processed, or all callbacks are disabled for all events while an event is being processed, then  
3220 when this callback routine returns, no more callbacks will be invoked for this instance of the event.

3221 Registering and unregistering a callback routine is a global operation and affects all threads, in a  
3222 multithreaded application. See Section 5.4.1 Multiple Callbacks.



### 3223 5.5.2. OpenACC Events During Event Processing

3224 OpenACC events may occur during event processing. This may be because of OpenACC API rou-  
3225 tine calls or OpenACC constructs being reached during event processing, or because of multiple host  
3226 threads executing asynchronously. Both the OpenACC runtime and the tool library must implement  
3227 the proper behavior.

### 3228 5.5.3. Multiple Host Threads

3229 Many programs that use OpenACC also use multiple host threads, such as programs using the  
3230 OpenMP API. The appearance of multiple host threads affects both the OpenACC runtime and the  
3231 tools library.

### 3232 Runtime Support for Multiple Threads

3233 The OpenACC runtime must be thread-safe, and the OpenACC runtime implementation of this  
3234 tools interface must also be thread-safe. All threads use the same set of callbacks for all events, so  
3235 registering a callback from one thread will cause all threads to execute that callback. This means that  
3236 managing the callback lists for each event must be protected from multiple simultaneous updates.  
3237 This includes adding a callback to the tail of the callback list for an event, removing a callback from  
3238 the list for an event, and incrementing or decrementing the *ref* count for a callback routine for an  
3239 event.

3240 In addition, one thread may register, unregister, enable, or disable a callback for an event while  
3241 another thread is processing the callback list for that event asynchronously. The exact behavior may  
3242 be dependent on the implementation, but some behaviors are expected and others are disallowed.  
3243 In the following examples, there are three callbacks, A, B, and C, registered for event E in that  
3244 order, where callbacks A and B are enabled and callback C is temporarily disabled. Thread T1 is  
3245 dynamically modifying the callbacks for event E while thread T2 is processing an instance of event  
3246 E.

- 3247 • Suppose thread T1 unregisters or disables callback A for event E. Thread T2 may or may not  
3248 invoke callback A for this event instance, but it must invoke callback B; if it invokes callback  
3249 A, that must precede the invocation of callback B.
- 3250 • Suppose thread T1 unregisters or disables callback B for event E. Thread T2 may or may not  
3251 invoke callback B for this event instance, but it must invoke callback A; if it invokes callback  
3252 B, that must follow the invocation of callback A.
- 3253 • Suppose thread T1 unregisters or disables callback A and then unregisters or disables callback  
3254 B for event E. Thread T2 may or may not invoke callback A and may or may not invoke  
3255 callback B for this event instance, but if it invokes both callbacks, it must invoke callback A  
3256 before it invokes callback B.
- 3257 • Suppose thread T1 unregisters or disables callback B and then unregisters or disables callback  
3258 A for event E. Thread T2 may or may not invoke callback A and may or may not invoke  
3259 callback B for this event instance, but if it invokes callback B, it must have invoked callback  
3260 A for this event instance.
- 3261 • Suppose thread T1 is registering a new callback D for event E. Thread T2 may or may not

3262        invoke callback D for this event instance, but it must invoke both callbacks A and B. If it  
3263        invokes callback D, that must follow the invocations of A and B.

- 3264        • Suppose thread T1 is enabling callback C for event E. Thread T2 may or may not invoke  
3265        callback C for this event instance, but it must invoke both callbacks A and B. If it invokes  
3266        callback C, that must follow the invocations of A and B.

3267        The `acc_prof_info` struct has a `thread_id` field, which the runtime must set to a unique  
3268        value for each host thread, though it need not be the same as the OpenMP threadnum value.

### 3269        **Library Support for Multiple Threads**

3270        The tool library must also be thread-safe. The callback routine will be invoked in the context of the  
3271        thread that reaches the event. The library may receive a callback from a thread T2 while it's still  
3272        processing a callback, from the same event type or from a different event type, from another thread  
3273        T1. The `acc_prof_info` struct has a `thread_id` field, which the runtime must set to a unique  
3274        value for each host thread.

3275        If the tool library uses dynamic callback registration and unregistration, or callback disabling and  
3276        enabling, recall that unregistering or disabling an event callback from one thread will unregister or  
3277        disable that callback for all threads, and registering or enabling an event callback from any thread  
3278        will register or enable it for all threads. If two or more threads register the same callback for the  
3279        same event, the behavior is the same as if one thread registered that callback multiple times; see  
3280        Section 5.4.1 Multiple Callbacks. The `acc_unregister` routine must be called as many times  
3281        as `acc_register` for that callback/event pair in order to totally unregister it. If two threads  
3282        register two different callback routines for the same event, unless the order of the registration calls  
3283        is guaranteed by some synchronization method, the order in which the runtime sees the registration  
3284        may differ for multiple runs, meaning the order in which the callbacks occur will differ as well.

## 6. Glossary

- 3286 Clear and consistent terminology is important in describing any programming model. We define  
3287 here the terms you must understand in order to make effective use of this document and the asso-  
3288 ciated programming model. In particular, some terms used in this specification conflict with their  
3289 usage in the base language specifications. When there is potential confusion, the term will appear  
3290 here.
- 3291 **Accelerator** – a device attached to a CPU and to which the CPU can offload data and compute  
3292 kernels to perform compute-intensive calculations.
- 3293 **Accelerator routine** – a C or C++ function or Fortran subprogram compiled for the accelerator  
3294 with the **routine** directive.
- 3295 **Accelerator thread** – a thread of execution that executes on the accelerator; a single vector lane of  
3296 a single worker of a single gang.
- 3297 **Aggregate datatype** – an array or composite datatype, or any non-scalar datatype. In Fortran, ag-  
3298 gregate datatypes include arrays and derived types. In C, aggregate datatypes include arrays, targets  
3299 of pointers, structs, and unions. In C++, aggregate datatypes include arrays, targets of pointers,  
3300 classes, structs, and unions.
- 3301 **Aggregate variables** – an array or composite variable, or a variable of any non-scalar datatype.
- 3302 **Async-argument** – an *async-argument* is a nonnegative scalar integer expression (*int* for C or C++,  
3303 *integer* for Fortran), or one of the special async values **acc\_async\_noval** or **acc\_async\_sync**.
- 3304 **Barrier** – a type of synchronization where all parallel execution units or threads must reach the  
3305 barrier before any execution unit or thread is allowed to proceed beyond the barrier; modeled after  
3306 the starting barrier on a horse race track.
- 3307 **Compute intensity** – for a given loop, region, or program unit, the ratio of the number of arithmetic  
3308 operations performed on computed data divided by the number of memory transfers required to  
3309 move that data between two levels of a memory hierarchy.
- 3310 **Construct** – a directive and the associated statement, loop, or structured block, if any.
- 3311 **Composite datatype** – a derived type in Fortran, or a **struct** or **union** type in C, or a **class**,  
3312 **struct**, or **union** type in C++. (This is different from the use of the term *composite data type* in  
3313 the C and C++ languages.)
- 3314 **Composite variable** – a variable of composite datatype. In Fortran, a composite variable must not  
3315 have allocatable or pointer attributes.
- 3316 **Compute construct** – a *parallel construct*, *kernels construct*, or *serial construct*.
- 3317 **Compute region** – a *parallel region*, *kernels region*, or *serial region*.
- 3318 **CUDA** – the CUDA environment from NVIDIA is a C-like programming environment used to  
3319 explicitly control and program an NVIDIA GPU.

- 3320 **Current device** – the device represented by the *acc-current-device-type-var* and *acc-current-device-*  
3321 *num-var* ICVs
- 3322 **Current device type** – the device type represented by the *acc-current-device-type-var* ICV
- 3323 **Data lifetime** – the lifetime of a data object in device memory, which may begin at the entry to  
3324 a data region, or at an **enter data** directive, or at a data API call such as **acc\_copyin** or  
3325 **acc\_create**, and which may end at the exit from a data region, or at an **exit data** directive,  
3326 or at a data API call such as **acc\_delete**, **acc\_copyout**, or **acc\_shutdown**, or at the end of  
3327 the program execution.
- 3328 **Data region** – a *region* defined by a **data** construct, or an implicit data region for a function or  
3329 subroutine containing OpenACC directives. Data constructs typically allocate device memory and  
3330 copy data from host to device memory upon entry, and copy data from device to local memory and  
3331 deallocate device memory upon exit. Data regions may contain other data regions and compute  
3332 regions.
- 3333 **Device** – a general reference to an accelerator or a multicore CPU.
- 3334 **Default asynchronous queue** – the asynchronous activity queue represented in the *acc-default-*  
3335 *async-var* ICV
- 3336 **Device memory** – memory attached to a device, logically and physically separate from the host  
3337 memory.
- 3338 **Device thread** – a thread of execution that executes on any device.
- 3339 **Directive** – in C or C++, a **#pragma**, or in Fortran, a specially formatted comment statement, that  
3340 is interpreted by a compiler to augment information about or specify the behavior of the program.
- 3341 **Discrete memory** – memory accessible from the local thread that is not accessible from the current  
3342 device, or memory accessible from the current device that is not accessible from the local thread.
- 3343 **DMA** – Direct Memory Access, a method to move data between physically separate memories;  
3344 this is typically performed by a DMA engine, separate from the host CPU, that can access the host  
3345 physical memory as well as an IO device or other physical memory.
- 3346 **GPU** – a Graphics Processing Unit; one type of accelerator.
- 3347 **GPGPU** – General Purpose computation on Graphics Processing Units.
- 3348 **Host** – the main CPU that in this context may have one or more attached accelerators. The host  
3349 CPU controls the program regions and data loaded into and executed on one or more devices.
- 3350 **Host thread** – a thread of execution that executes on the host.
- 3351 **Implicit data region** – the data region that is implicitly defined for a Fortran subprogram or C  
3352 function. A call to a subprogram or function enters the implicit data region, and a return from the  
3353 subprogram or function exits the implicit data region.
- 3354 **Kernel** – a nested loop executed in parallel by the accelerator. Typically the loops are divided into  
3355 a parallel domain, and the body of the loop becomes the body of the kernel.
- 3356 **Kernels region** – a *region* defined by a **kernels** construct. A kernels region is a structured block  
3357 which is compiled for the accelerator. The code in the kernels region will be divided by the compiler  
3358 into a sequence of kernels; typically each loop nest will become a single kernel. A kernels region  
3359 may require space in device memory to be allocated and data to be copied from local memory to

- 3360 device memory upon region entry, and data to be copied from device memory to local memory and  
3361 space in device memory to be deallocated upon exit.
- 3362 **Level of parallelism** – The possible levels of parallelism in OpenACC are gang, worker, vector,  
3363 and sequential. One or more of gang, worker, and vector parallelism may be specified on a loop  
3364 construct. Sequential execution corresponds to no parallelism. The **gang**, **worker**, **vector**, and  
3365 **seq** clauses specify the level of parallelism for a loop.
- 3366 **Local device** – the device where the *local thread* executes.
- 3367 **Local memory** – the memory associated with the *local thread*.
- 3368 **Local thread** – the host thread or the accelerator thread that executes an OpenACC directive or  
3369 construct.
- 3370 **Loop trip count** – the number of times a particular loop executes.
- 3371 **MIMD** – a method of parallel execution (Multiple Instruction, Multiple Data) where different exe-  
3372 cution units or threads execute different instruction streams asynchronously with each other.
- 3373 **OpenCL** – short for Open Compute Language, a developing, portable standard C-like programming  
3374 environment that enables low-level general-purpose programming on GPUs and other accelerators.
- 3375 **Orphaned loop construct** - a **loop** construct that is not lexically contained in any compute con-  
3376 struct, that is, that has no parent compute construct.
- 3377 **Parallel region** – a *region* defined by a **parallel** construct. A parallel region is a structured block  
3378 which is compiled for the accelerator. A parallel region typically contains one or more work-sharing  
3379 loops. A parallel region may require space in device memory to be allocated and data to be copied  
3380 from local memory to device memory upon region entry, and data to be copied from device memory  
3381 to local memory and space in device memory to be deallocated upon exit.
- 3382 **Parent compute construct** – for a **loop** construct, the **parallel**, **kernels**, or **serial** con-  
3383 struct that lexically contains the **loop** construct and is the innermost compute construct that con-  
3384 tains that **loop** construct, if any.
- 3385 **Present data** – data for which the sum of the structured and dynamic reference counters is greater  
3386 than zero.
- 3387 **Private data** – with respect to an iterative loop, data which is used only during a particular loop  
3388 iteration. With respect to a more general region of code, data which is used within the region but is  
3389 not initialized prior to the region and is re-initialized prior to any use after the region.
- 3390 **Procedure** – in C or C++, a function in the program; in Fortran, a subroutine or function.
- 3391 **Region** – all the code encountered during an instance of execution of a construct. A region includes  
3392 any code in called routines, and may be thought of as the dynamic extent of a construct. This may  
3393 be a *parallel region*, *kernels region*, *serial region*, *data region* or *implicit data region*.
- 3394 **Scalar** – a variable of scalar datatype. In Fortran, scalars must not have allocatable or pointer  
3395 attributes.
- 3396 **Scalar datatype** – an intrinsic or built-in datatype that is not an array or aggregate datatype. In For-  
3397 tran, scalar datatypes are integer, real, double precision, complex, or logical. In C, scalar datatypes  
3398 are char (signed or unsigned), int (signed or unsigned, with optional short, long or long long at-  
3399 tribute), enum, float, double, long double, `_Complex` (with optional float or long attribute), or any  
3400 pointer datatype. In C++, scalar datatypes are char (signed or unsigned), `wchar_t`, int (signed or

3401 unsigned, with optional short, long or long long attribute), enum, bool, float, double, long double,  
3402 or any pointer datatype. Not all implementations or targets will support all of these datatypes.

3403 **Serial region** – a *region* defined by a **serial** construct. A serial region is a structured block which  
3404 is compiled for the accelerator. A serial region contains code that is executed by one vector lane of  
3405 one worker in one gang. A serial region may require space in device memory to be allocated and  
3406 data to be copied from local memory to device memory upon region entry, and data to be copied  
3407 from device memory to local memory and space in device memory to be deallocated upon exit.

3408 **Shared memory** – memory that is accessible from both the local thread and the current device.

3409 **SIMD** – A method of parallel execution (single-instruction, multiple-data) where the same instruc-  
3410 tion is applied to multiple data elements simultaneously.

3411 **SIMD operation** – a *vector operation* implemented with SIMD instructions.

3412 **Structured block** – in C or C++, an executable statement, possibly compound, with a single entry  
3413 at the top and a single exit at the bottom. In Fortran, a block of executable statements with a single  
3414 entry at the top and a single exit at the bottom.

3415 **Thread** – On a host CPU, a thread is defined by a program counter and stack location; several host  
3416 threads may comprise a process and share host memory. On an accelerator, a thread is any one  
3417 vector lane of one worker of one gang.

3418 **var** – the name of a variable (scalar, array, or composite variable), or a subarray specification, or an  
3419 array element, or a composite variable member, or the name of a Fortran common block between  
3420 slashes.

3421 **Vector operation** – a single operation or sequence of operations applied uniformly to each element  
3422 of an array.

3423 **Visible device copy** – a copy of a variable, array, or subarray allocated in device memory that is  
3424 visible to the program unit being compiled.

## 3425 **A. Recommendations for Implementors**

3426 This section gives recommendations for standard names and extensions to use for implementations  
3427 for specific targets and target platforms, to promote portability across such implementations, and  
3428 recommended options that programmers find useful. While this appendix is not part of the Open-  
3429 ACC specification, implementations that provide the functionality specified herein are strongly rec-  
3430 ommended to use the names in this section. The first subsection describes devices, such as NVIDIA  
3431 GPUs. The second subsection describes additional API routines for target platforms, such as CUDA  
3432 and OpenCL. The third subsection lists several recommended options for implementations.

### 3433 **A.1. Target Devices**

#### 3434 **A.1.1. NVIDIA GPU Targets**

3435 This section gives recommendations for implementations that target NVIDIA GPU devices.

#### 3436 **Accelerator Device Type**

3437 These implementations should use the name **acc\_device\_nvidia** for the **acc\_device\_t**  
3438 type or return values from OpenACC Runtime API routines.

#### 3439 **ACC\_DEVICE\_TYPE**

3440 An implementation should use the case-insensitive name **nvidia** for the environment variable  
3441 **ACC\_DEVICE\_TYPE**.

#### 3442 **device.type clause argument**

3443 An implementation should use the case-insensitive name **nvidia** as the argument to the **device.type**  
3444 clause.

#### 3445 **A.1.2. AMD GPU Targets**

3446 This section gives recommendations for implementations that target AMD GPUs.

### 3447 **Accelerator Device Type**

3448 These implementations should use the name **acc\_device\_radeon** for the **acc\_device\_t**  
3449 type or return values from OpenACC Runtime API routines.

### 3450 **ACC\_DEVICE\_TYPE**

3451 These implementations should use the case-insensitive name **radeon** for the environment variable  
3452 **ACC\_DEVICE\_TYPE**.

### 3453 **device\_type clause argument**

3454 An implementation should use the case-insensitive name **radeon** as the argument to the **device\_type**  
3455 clause.

## 3456 **A.1.3. Multicore Host CPU Target**

3457 This section gives recommendations for implementations that target the multicore host CPU.

### 3458 **Accelerator Device Type**

3459 These implementations should use the name **acc\_device\_host** for the **acc\_device\_t** type  
3460 or return values from OpenACC Runtime API routines.

### 3461 **ACC\_DEVICE\_TYPE**

3462 These implementations should use the case-insensitive name **host** for the environment variable  
3463 **ACC\_DEVICE\_TYPE**.

### 3464 **device\_type clause argument**

3465 An implementation should use the case-insensitive name **host** as the argument to the **device\_type**  
3466 clause.

## 3467 **A.2. API Routines for Target Platforms**

3468 These runtime routines allow access to the interface between the OpenACC runtime API and the  
3469 underlying target platform. An implementation may not implement all these routines, but if it  
3470 provides this functionality, it should use these function names.



### 3471 **A.2.1. NVIDIA CUDA Platform**

3472 This section gives runtime API routines for implementations that target the NVIDIA CUDA Run-  
3473 time or Driver API.

#### 3474 **acc\_get\_current\_cuda\_device**

3475 **Summary** The `acc_get_current_cuda_device` routine returns the NVIDIA CUDA de-  
3476 vice handle for the current device.

#### 3477 **Format**

C or C++:

```
void* acc_get_current_cuda_device ();
```

#### 3478 **acc\_get\_current\_cuda\_context**

3479 **Summary** The `acc_get_current_cuda_context` routine returns the NVIDIA CUDA  
3480 context handle in use for the current device.

#### 3481 **Format**

C or C++:

```
void* acc_get_current_cuda_context ();
```

#### 3482 **acc\_get\_cuda\_stream**

3483 **Summary** The `acc_get_cuda_stream` routine returns the NVIDIA CUDA stream handle  
3484 in use for the current device for the specified async value.

#### 3485 **Format**

C or C++:

```
void* acc_get_cuda_stream ( int async );
```

#### 3486 **acc\_set\_cuda\_stream**

3487 **Summary** The `acc_set_cuda_stream` routine sets the NVIDIA CUDA stream handle the  
3488 current device for the specified async value.

#### 3489 **Format**

C or C++:

```
int acc_set_cuda_stream ( int async, void* stream );
```

## 3490 A.2.2. OpenCL Target Platform

3491 This section gives runtime API routines for implementations that target the OpenCL API on any  
3492 device.

### 3493 **acc\_get\_current\_opengl\_device**

3494 **Summary** The **acc\_get\_current\_opengl\_device** routine returns the OpenCL device  
3495 handle for the current device.

#### 3496 **Format**

C or C++:

```
void* acc_get_current_opengl_device ();
```

### 3497 **acc\_get\_current\_opengl\_context**

3498 **Summary** The **acc\_get\_current\_opengl\_context** routine returns the OpenCL context  
3499 handle in use for the current device.

#### 3500 **Format**

C or C++:

```
void* acc_get_current_opengl_context ();
```

### 3501 **acc\_get\_opengl\_queue**

3502 **Summary** The **acc\_get\_opengl\_queue** routine returns the OpenCL command queue han-  
3503 dle in use for the current device for the specified async value.

#### 3504 **Format**

C or C++:

```
cl_command_queue acc_get_opengl_queue ( int async );
```

### 3505 **acc\_set\_opengl\_queue**

3506 **Summary** The **acc\_set\_opengl\_queue** routine returns the OpenCL command queue han-  
3507 dle in use for the current device for the specified async value.

#### 3508 **Format**

C or C++:

```
void acc_set_opengl_queue ( int async, cl_command_queue cmdqueue );
```

### 3509 A.3. Recommended Options

3510 The following options are recommended for implementations; for instance, these may be imple-  
3511 mented as command-line options to a compiler or settings in an IDE.

#### 3512 A.3.1. C Pointer in Present clause

3513 This revision of OpenACC clarifies the construct:

```
void test(int n ){  
    float* p;  
    ...  
    #pragma acc data present (p)  
    {  
        // code here...  
    }
```

3514 This example tests whether the pointer **p** itself is present in the current device memory. Implemen-  
3515 tations before this revision commonly implemented this by testing whether the pointer target **p[0]**  
3516 was present in the current device memory, and this appears in many programs assuming such. Until  
3517 such programs are modified to comply with this revision, an option to implement **present (p)** as  
3518 **present (p[0])** for C pointers may be helpful to users.

#### 3519 A.3.2. Autoscopying

3520 If an implementation implements autoscopying to automatically determine variables that are private  
3521 to a compute region or to a loop, or to recognize reductions in a compute region or a loop, an option  
3522 to print a message telling what variables were affected by the analysis would be helpful to users. An  
3523 option to disable the autoscopying analysis would be helpful to promote program portability across  
3524 implementations.



# Index

- 3525 **\_OPENACC**, 12, 14–16, 20, 107
- 3526 **acc-current-device-num-var**, 20
- 3527 **acc-current-device-type-var**, 20
- 3528 **acc-default-async-var**, 20, 71
- 3529 **acc\_async\_noval**, 12, 71
- 3530 **acc\_async\_sync**, 12, 71
- 3531 **acc\_device\_host**, 128
- 3532 **ACC\_DEVICE\_NUM**, 21, 99
- 3533 **acc\_device\_nvidia**, 127
- 3534 **acc\_device\_radeon**, 128
- 3535 **ACC\_DEVICE\_TYPE**, 21, 99, 127, 128
- 3536 **ACC\_PROFLIB**, 99
- 3537 **action**
  - 3538 **attach**, 35, 40
  - 3539 **copyin**, 39
  - 3540 **copyout**, 39
  - 3541 **create**, 39
  - 3542 **delete**, 39
  - 3543 **detach**, 35, 40
  - 3544 **immediate**, 40
  - 3545 **present decrement**, 38
  - 3546 **present increment**, 38
- 3547 **AMD GPU target**, 127
- 3548 **async** clause, 34, 66, 71
- 3549 **async queue**, 9
- 3550 *async-argument*, 71
- 3551 **asynchronous execution**, 9, 71
- 3552 **atomic** construct, 13, 54
- 3553 **attach** action, 35, 40
- 3554 **attach** clause, 45
- 3555 **attachment counter**, 35
- 3556 **auto** clause, 13, 50
- 3557 **autoscopying**, 131
- 3558 **barrier synchronization**, 9, 23, 24, 26, 123
- 3559 **bind** clause, 69
- 3560 **cache** directive, 52
- 3561 **capture** clause, 58
- 3562 **collapse** clause, 48
- 3563 **common block**, 36, 59, 61, 70
- 3564 **compute construct**, 123
- 3565 **compute region**, 123
- 3566 **construct**, 123
  - 3567 **atomic**, 54
  - 3568 **compute**, 123
  - 3569 **data**, 32, 36
  - 3570 **host\_data**, 45
  - 3571 **kernels**, 23, 24, 36
  - 3572 **kernels loop**, 53
  - 3573 **parallel**, 22, 36
  - 3574 **parallel loop**, 53
  - 3575 **serial**, 25, 36
  - 3576 **serial loop**, 53
- 3577 **copy** clause, 42
- 3578 **copyin** action, 39
- 3579 **copyin** clause, 42
- 3580 **copyout** action, 39
- 3581 **copyout** clause, 43
- 3582 **create** action, 39
- 3583 **create** clause, 43, 61
- 3584 **CUDA**, 9, 10, 123, 127, 129
- 3585 **data** attribute
  - 3586 **explicitly determined**, 30
  - 3587 **implicitly determined**, 30
  - 3588 **predetermined**, 30
- 3589 **data** clause, 36
- 3590 **data** construct, 32, 36
- 3591 **data** lifetime, 124
- 3592 **data** region, 31, 124
  - 3593 **implicit**, 31
- 3594 **declare** directive, 13, 59
- 3595 **default** clause, 29
- 3596 **default (none)** clause, 12, 23, 24, 26
- 3597 **default(present)**, 23, 24, 26
- 3598 **delete** action, 39
- 3599 **delete** clause, 44
- 3600 **detach** action, 35, 40
  - 3601 **immediate**, 40

- 3602 **detach** clause, 45
- 3603 **device** clause, 66
- 3604 **device\_resident** clause, 60
- 3605 **device\_type** clause, 13, 21, 36, 127, 128
- 3606 **deviceptr** clause, 36, 41
- 3607 direct memory access, 9, 124
- 3608 DMA, 9, 124
- 3609 **enter data** directive, 33, 36
- 3610 environment variable
- 3611     **\_OPENACC**, 20
- 3612     **ACC\_DEVICE\_NUM**, 21, 99
- 3613     **ACC\_DEVICE\_TYPE**, 21, 99, 127, 128
- 3614     **ACC\_PROFLIB**, 99
- 3615 **exit data** directive, 33, 36
- 3616 explicitly determined data attribute, 30
- 3617 **firstprivate** clause, 23, 26, 28
- 3618 gang, 23, 26
- 3619 **gang** clause, 48, 68
- 3620 gang parallelism, 8
- 3621 *gang-arg*, 47
- 3622 gang-partitioned mode, 8
- 3623 gang-redundant mode, 8, 23, 26
- 3624 GP mode, 8
- 3625 GR mode, 8
- 3626 **host**, 128
- 3627 **host** clause, 13, 66
- 3628 **host\_data** construct, 45
- 3629 ICV, 20
- 3630 **if** clause, 32, 34, 66
- 3631 immediate detach action, 40
- 3632 implicit data region, 31
- 3633 implicitly determined data attribute, 30
- 3634 **independent** clause, 51
- 3635 **init** directive, 62
- 3636 internal control variable, 20
- 3637 **kernels** construct, 23, 24, 36
- 3638 **kernels loop** construct, 53
- 3639 level of parallelism, 8, 125
- 3640 **link** clause, 13, 36, 61
- 3641 local device, 9
- 3642 local memory, 9
- 3643 local thread, 9
- 3644 **loop** construct, 47
- 3645     orphaned, 48
- 3646 **no\\_create** clause, 44
- 3647 **nohost** clause, 70
- 3648 **num\_gangs** clause, 27
- 3649 **num\_workers** clause, 27
- 3650 **nvidia**, 127
- 3651 NVIDIA GPU target, 127
- 3652 OpenCL, 9, 10, 125, 127, 130
- 3653 orphaned **loop** construct, 48
- 3654 **parallel** construct, 22, 36
- 3655 **parallel loop** construct, 53
- 3656 parallelism
- 3657     level, 8, 125
- 3658 parent compute construct, 48
- 3659 predetermined data attribute, 30
- 3660 **present** clause, 36, 41
- 3661 present decrement action, 38
- 3662 present increment action, 38
- 3663 **private** clause, 28, 51
- 3664 **radeon**, 128
- 3665 **read** clause, 58
- 3666 **reduction** clause, 28, 51
- 3667 reference counter, 35
- 3668 region
- 3669     compute, 123
- 3670     data, 31, 124
- 3671     implicit data, 31
- 3672 **routine** directive, 13, 68
- 3673 **self** clause, 13, 66
- 3674 sentinel, 19
- 3675 **seq** clause, 50, 69
- 3676 **serial** construct, 25, 36
- 3677 **serial loop** construct, 53
- 3678 **shutdown** directive, 63
- 3679 *size-expr*, 47
- 3680 thread, 126
- 3681 **tile** clause, 13, 50
- 3682 **update** clause, 58
- 3683 **update** directive, 65
- 3684 **use\_device** clause, 46
- 3685 **vector** clause, 49, 69
- 3686 vector lane, 23

- 3687 vector parallelism, 8
- 3688 vector-partitioned mode, 8
- 3689 vector-single mode, 8
- 3690 **vector\_length** clause, 28
- 3691 visible device copy, 126
- 3692 VP mode, 8
- 3693 VS mode, 8
  
- 3694 **wait** clause, 34, 66, 72
- 3695 **wait** directive, 72
- 3696 worker, 23, 26
- 3697 **worker** clause, 49, 69
- 3698 worker parallelism, 8
- 3699 worker-partitioned mode, 8
- 3700 worker-single mode, 8
- 3701 WP mode, 8
- 3702 WS mode, 8