



The OpenMP Architecture Review Board

OpenMP Technical Report 5: Memory Management Support for OpenMP 5.0

This Technical Report augments the OpenMP TR 4 document with language features for managing memory on systems with heterogeneous memories.

EDITORS

Alejandro Duran (alejandro.duran@intel.com)

Christian Terboven (terboven@itc.rwth-aachen.de)

OTHER AUTHORS

Jonathan Beard (ARM)

Ravi Narayanaswamy (Intel)

Bronis de Supinski (LLNL)

John Pennycook (Intel)

Deepak Eachempati (Cray)

Alejandro Rico (ARM)

Alexandre Eichenberger (IBM)

Jeff Sandoval (Cray)

Ian Karlin (LLNL)

Tom Scogland (LLNL)

Kelvin Li (IBM)

Jason Sewall (Intel)

Stephen Olivier (SNL)

Xinmin Tian (Intel)

and other members of the Affinity Subcommittee

January 5th, 2017

Expires January 4th, 2018

We actively solicit comments. Please provide feedback on this document either to the Editor directly or in the OpenMP Forum at openmp.org

End of Public Comment Period: March 5th, 2017

OpenMP Architecture Review Board

www.openmp.org info@openmp.org

OpenMP ARB – Ravi S. Rao, c/o Intel Corporation, 1300 S MoPac Express Way, Austin, TX 78746, USA

This technical report describes possible future directions or extensions to the OpenMP Specification.

The goal of this technical report is to build more widespread existing practice for an expanded OpenMP. It gives advice on extensions or future directions to those vendors who wish to provide them possibly for trial implementation, allows OpenMP to gather early feedback, support timing and scheduling differences between official OpenMP releases, and offers a preview to users of the future directions of OpenMP with the provision stated in the next paragraph.

This technical report is non-normative. Some of the components in this technical report may be considered for standardization in a future version of OpenMP, but they are not currently part of any OpenMP Specification. Some of the components in this technical report may never be standardized, others may be standardized in a substantially changed form, or it may be standardized as is in its entirety.

Memory Management support for OpenMP

The OpenMP Affinity Subcommittee

1 Motivation and Background

System performance is often dependent on memory performance. Over the past decades the bandwidth of the standard memory technology (DRAM) has scaled slower than the increase in CPU computational throughput. System builders traditionally addressed this problem by adding more memory channels to maintain system balance. However, recently, bandwidth and capacity are scaling slower than compute and vendors have not been able to maintain system balance with DRAM-only solutions. To address this problem, emerging systems feature multiple types of memories with different optimization points. Examples are systems that combine off-package DRAM with higher bandwidth technologies integrated on package to increase memory bandwidth, non-volatile high-density memories to increase capacity, and on-chip scratchpad memories with low-latency access.

Compute systems with such a tiered memory solution present a unique challenge to programmers. With the fastest resources typically having limited capacity, placement choices present performance tradeoffs in applications. Also, traditional first touch placement strategies used in Linux do not allow users to differentiate among memories with different properties. Vendors provide their own programming approaches to differentiate different memories, e.g., CUDA and Memkind, but these approaches and the alternative low-level programming approaches are non-portable. In response, to enable portability across platforms, the OpenMP committee is developing a more consistent and portable interface for memory placement in tiered memory systems.

The proposal in this document is designed to abstract the myriad of choices from the user. The goal is to enable portability, while providing the user with enough control to allow a runtime to manage allocations for user-defined properties such as latency, bandwidth and capacity. We aim to use properties and traits rather than the specific memory types of today to help future proof the interface against emerging and changing technology trends.

This document represents current directions being discussed within the OpenMP Affinity Subcommittee and is designed to engage the community, solicit feedback and reflect the current thoughts of the committee on this topic. The proposal is a start of a larger document that will include controls to cover additional memory types and features, such as persistent

memory and constant memory. This document is not a promise that the interface will be adopted into the specification. Instead, it represents the subcommittee's best estimate of a portion of an interface that will be adopted, assuming that the OpenMP community agrees that the interface can be extended to fully support the range of architectures of interest.

2 High-level overview

A platform-agnostic integration of memory management support into OpenMP is necessary to avoid the separation of code paths for different platforms and also different kinds of memory within each platform. As a de-facto standard, OpenMP has to support all current kinds of memory and has to be capable of supporting future memory kinds and platform configuration without significant changes to both the specification and any code using the OpenMP memory management. This is achieved by introducing the following new concepts into the OpenMP API:

- **Memory spaces and allocators:** A *memory space* refers to a memory resource available in the system at the time the OpenMP program is executed. Each space has certain characteristics depending on the kind of the physical memory and the current state of the system. An *allocator* is an object that allocates (and frees) memory from an associated memory space.
- **Memory allocation API:** The `omp_alloc()` and `omp_free()` API routines are provided for C/C++ to allocate and deallocate memory using an allocator.
- **`allocate` directive and clause:** The new `allocate` directive and clause allow the allocation of variables without the explicit use of the aforementioned API, and can be used in both Fortran and C/C++. They support several modifiers to influence their behavior.

In order to work with memory spaces and allocators, an API is provided to manage (i.e., create and destroy) both types of objects. The programmer must explicitly use this API to enable the use of memory types other than the default type with OpenMP.

The mixture of run-time and compile-time functionality is necessary to handle the different types of memory allocations, namely a `malloc()`-like interface for dynamic (heap) allocations in C/C++ and directives for static and stack allocations in both Fortran and C/C++. A mixture of runtime and compiler support is also necessary to support certain kinds of memory that need special (machine) instructions to access or modify data.

With respect to future architectural developments, it must be assumed that hardware will develop at a faster rate than the OpenMP specification can match. In consequence, the options to express certain memory properties are not tied to current systems. Instead, the options aim to be broadly applicable by referring to certain characteristics of memory resources, and they are intended to be extended by vendors with the introduction of additional traits.

2.1 Memory spaces and allocators

A *memory space* represents a storage resource that is available in the system. For example, almost all contemporary HPC systems contain a DDR-based main memory, which could be the only available memory space. Additional new memory types include those with enhanced performance (e.g., high-bandwidth memory) or functionality (e.g., non-volatile memory).

Both could be additional memory spaces in a single system, and numerous combinations are possible.

A memory space is represented by the `omp_memspace_t` C/C++ datatype (`omp_memspace_kind` in Fortran). Before first use, it has to be initialized via the corresponding initialization function `omp_init_memspace`, which accepts a set of memory traits (see next paragraph) as the argument. The instance of a memory space is itself passed as an argument in the construction of an allocator. After last use, the memory space must be destroyed via `omp_destroy_memspace`.

Memory traits describe the characteristics of memory spaces and as such allow for queries, identification and description of the different memory spaces of a system. This proposal contains a base set of memory traits described below, others may be added in the future or as vendor-specific extensions. Memory traits can either be prescriptive, meaning an exact match is required, or descriptive, meaning the runtime is requested to select the optimal type of memory based on the requested properties.

Prescriptive traits include the location of memory (with possible values core, socket or device), a certain optimization characteristic of the underlying memory technology (with possible values bandwidth or latency or capacity), and support for certain page sizes or read/write permission. Descriptive traits include the relative distance relative to the task performing the request (with possible values near or far) and the relative bandwidth and latency of the memory space with respect to other memories in the system (with possible values highest and lowest).

A memory trait is represented by the `omp_memtrait_t` datatype and support for sets of memory traits is represented by `omp_memtrait_set_t` in C/C++, with corresponding Fortran types/kinds. The `omp_init_memtrait_set` API routine is available to construct a memory trait set from a given list of memory traits. The trait set is used as an argument to `omp_init_memspace`, with `omp_default_memtraits` representing the default memory as selected by the runtime. Traits to request a minimum total capacity and available capacity are also available. Associated routines include `omp_destroy_memtrait_set`, to destroy the memtrait set, `omp_add_memtraits`, to add a memory trait to a memory traits set, and `omp_merge_memtraits`, to merge two memory traits sets.

An *allocator* is an object performing allocations of contiguous memory chunks from a given memory space. *Allocator traits* can be employed to customize the behavior of an allocator. This includes the behavior in case the allocation is not successful – the standard behavior in case of failure is to fall back to the default memory, based on the `omp_default_memtraits` specified at initialization of the memory space. On many systems that would be DDR main memory. Further allocator traits specify the thread model (with possible values shared or exclusive) and the options to specify alignment and the request for pinned memory.

An allocator is represented by the `omp_allocator_t` datatype (`omp_allocator_kind` in Fortran). Before first use, it has to be initialized via the corresponding initialization function `omp_init_allocator`, which accepts a memory space and a set of allocator traits as arguments. API routines for the management of allocator traits are similar to those for memory traits. After last use, the allocator has to be destroyed via `omp_destroy_allocator`.

2.2 Memory allocation API for C/C++

Two new API routines are provided to allocate and deallocate memory using an allocator in C/C++. Allocations are performed with the `omp_alloc` routine, which takes the requested size as the first argument and an OpenMP allocator as the second argument and returns a pointer to the allocated memory. The additional `omp_alloc_safe_align` routine requests an aligned allocation. Similarly, the `omp_free` routine frees memory and also takes an OpenMP allocator as the second argument. When memory of a given size is requested, memory of at least that size is allocated, and it must be freed with the corresponding function using the corresponding allocator.

The separation of the API and the allocators allows the programmer to write portable code because only the allocator definition must be modified when the code is changed to target a different kind of memory on a different platform, while all the individual allocations in the code can remain unmodified.

2.3 Allocate directive and clause

The new `allocate` directive enables the programmer to influence the allocation of variables without the explicit use of the aforementioned API. It also integrates the memory management concept with the other directives and constructs in the OpenMP API. The effect of using the `allocate` directive is that for all variables in the list the storage location is determined by the application of the given allocator object. The allocator can be specified via the `allocator` clause. If no allocator is given, an implicit allocator is constructed from the memory and allocator traits specified with the directive via the `memtraits` and `alloctraits` clauses, taking as arguments the corresponding trait sets as discussed above.

In Fortran, the `allocate` directive provides in addition to the semantics described above the ability to use the allocator functionality with variables declared as `ALLOCATABLE`. That means it ensures the following Fortran `ALLOCATE` statement is performed with the OpenMP allocator specified either explicitly or constructed implicitly from the provided trait sets.

For directives supporting the new `allocate` clause, it specifies the allocation and memory traits of the storage used for private variables of a directive.

2.4 Default allocator

The new def-allocator-var ICV determines the allocator to be used by allocation routines, directives and clauses when an allocator is not specified by the user. The new corresponding API routines `omp_get_default_allocator` and `omp_set_default_allocator` are introduced, along with the new environment variable `OMP_ALLOCATOR`.

3 Changes to the OpenMP specification

In this section we present the necessary changes to be enacted to OpenMP TR4 document to enable our proposal. The new text that would be added is marked in blue and to simplify the presentation of the changes pages where the only changes are cross-references are not showed in this document.

3.1 Changes to Chapter 1

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

7 1.4.2 Device Data Environments

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

13 When an original variable is mapped to a device data environment and the associated
14 corresponding variable is not present in the device data environment, a new corresponding variable
15 (of the same type and size as the original variable) is created in the device data environment. The
16 initial value of the new corresponding variable is determined from the clauses and the data
17 environment of the encountering thread.

18 The corresponding variable in the device data environment may share storage with the original
19 variable. Writes to the corresponding variable may alter the value of the original variable. The
20 impact of this on memory consistency is discussed in Section 1.4.5 on page 21. When a task
21 executes in the context of a device data environment, references to the original variable refer to the
22 corresponding variable in the device data environment.

23 The relationship between the value of the original variable and the initial or final value of the
24 corresponding variable depends on the *map-type*. Details of this issue, as well as other issues with
25 mapping a variable, are provided in Section 2.17.6.1 on page 256.

26 The original variable in a data environment and the corresponding variable(s) in one or more device
27 data environments may share storage. Without intervening synchronization data races can occur.

28 1.4.3 Memory management

29 The host device, and target devices that an implementation may support, have attached
30 storage resources where program variables are stored. These resources can be of
31 different kinds and of different traits. A memory space in an OpenMP program represents
32 one of these resources. Memory spaces have different traits that define them and a single

1 resource may be exposed as multiple memory spaces with different traits. In any device at
2 least one memory space is guaranteed to exist.

3 An OpenMP program can use an allocator to allocate storage for its variables. Allocators
4 are associated with a memory space when created and use storage in that memory space
5 to allocate variables. Allocators are also used to deallocate variables and free the storage
6 in the memory space. When an OpenMP allocator is not used variables can be allocated
7 in any memory space. The behavior of a memory management construct, modifier or API
8 is unspecified if the variable that is applied to was not allocated with an OpenMP allocator.

9 1.4.4 The Flush Operation

10 The memory model has relaxed-consistency because a thread's temporary view of memory is not
11 required to be consistent with memory at all times. A value written to a variable can remain in the
12 thread's temporary view until it is forced to memory at a later time. Likewise, a read from a variable
13 may retrieve the value from the thread's temporary view, unless it is forced to read from memory.
14 The OpenMP flush operation enforces consistency between the temporary view and memory.

15 The flush operation is applied to a set of variables called the *flush-set*. The flush operation restricts
16 reordering of memory operations that an implementation might otherwise do. Implementations
17 must not reorder the code for a memory operation for a given variable, or the code for a flush
18 operation for the variable, with respect to a flush operation that refers to the same variable.

19 If a thread has performed a write to its temporary view of a shared variable since its last flush of
20 that variable, then when it executes another flush of the variable, the flush does not complete until
21 the value of the variable has been written to the variable in memory. If a thread performs multiple
22 writes to the same variable between two flushes of that variable, the flush ensures that the value of
23 the last write is written to the variable in memory. A flush of a variable executed by a thread also
24 causes its temporary view of the variable to be discarded, so that if its next memory operation for
25 that variable is a read, then the thread will read from memory when it may again capture the value
26 in the temporary view. When a thread executes a flush, no later memory operation by that thread for
27 a variable involved in that flush is allowed to start until the flush completes. The completion of a
28 flush of a set of variables executed by a thread is defined as the point at which all writes to those
29 variables performed by the thread before the flush are visible in memory to all other threads and
30 that thread's temporary view of all variables involved is discarded.

31 The flush operation provides a guarantee of consistency between a thread's temporary view and
32 memory. Therefore, the flush operation can be used to guarantee that a value written to a variable
33 by one thread may be read by a second thread. To accomplish this, the programmer must ensure
34 that the second thread has not written to the variable since its last flush of the variable, and that the
35 following sequence of events happens in the specified order:

36 1. The value is written to the variable by the first thread.

3.2 Changes to Chapter 2

Fortran

1 A *list item* is a variable, array section or common block name (enclosed in slashes). An *extended*
2 *list item* is a *list item* or a procedure name. A *locator list item* is a *list item*.

3 When a named common block appears in a *list*, it has the same meaning as if every explicit member
4 of the common block appeared in the list. An explicit member of a common block is a variable that
5 is named in a **COMMON** statement that specifies the common block name and is declared in the same
6 scoping unit in which the clause appears.

7 Although variables in common blocks can be accessed by use association or host association,
8 common block names cannot. As a result, a common block name specified in a data-sharing
9 attribute, a data copying or a data-mapping attribute clause must be declared to be a common block
10 in the same scoping unit in which the clause appears.

Fortran

11 For all base languages, a *list item* or an *extended list item* is subject to the restrictions specified in
12 Section 2.4 on page 48 and in each of the sections describing clauses and directives for which the
13 *list* or *extended-list* appears.

14 The clauses of the **allocate** directive accept a key-value list. A key-value list is a
15 comma-separated list of key-value pairs. A key-value pair has the form of key=value. The
16 allowed keys and values depend on each clause.

3.2.1 Changes to ICVs descriptions

1 • *bind-var* - controls the binding of OpenMP threads to places. When binding is requested, the
2 variable indicates that the execution environment is advised not to move threads between places.
3 The variable can also provide default thread affinity policies. There is one copy of this ICV per
4 data environment.

5 The following ICVs store values that affect the operation of loop regions.

6 • *run-sched-var* - controls the schedule that the **runtime** schedule clause uses for loop regions.
7 There is one copy of this ICV per data environment.

8 • *def-sched-var* - controls the implementation defined default scheduling of loop regions. There is
9 one copy of this ICV per device.

10 The following ICVs store values that affect program execution.

11 • *stacksize-var* - controls the stack size for threads that the OpenMP implementation creates. There
12 is one copy of this ICV per device.

13 • *wait-policy-var* - controls the desired behavior of waiting threads. There is one copy of this ICV
14 per device.

15 • *cancel-var* - controls the desired behavior of the **cancel** construct and cancellation points.
16 There is one copy of this ICV for the whole program.

17 • *default-device-var* - controls the default target device. There is one copy of this ICV per data
18 environment.

19 • *max-task-priority-var* - controls the maximum priority value that can be specified in the
20 **priority** clause of the **task** construct. There is one copy of this ICV for the whole program.

21 The following ICVs store values that affect the operation of the tool interface.

22 • *tool-var* - determines whether an OpenMP implementation will try to register a tool. There is
23 one copy of this ICV for the whole program.

24 • *tool-libraries-var* - specifies a list of absolute paths to tool libraries for OpenMP devices. There
25 is one copy of this ICV for the whole program.

26 The following ICVs store values that affect default memory allocation.

27 • *def-allocator-var* - determines the allocator to be used by allocation routines, directives
28 and clauses when an allocator is not specified by the user.

29 2.3.2 ICV Initialization

30 Table 2.1 shows the ICVs, associated environment variables, and initial values.

TABLE 2.1: ICV Initial Values

ICV	Environment Variable	Initial value
<i>dyn-var</i>	OMP_DYNAMIC	See description below
<i>nest-var</i>	OMP_NESTED	<i>false</i>
<i>nthreads-var</i>	OMP_NUM_THREADS	Implementation defined
<i>run-sched-var</i>	OMP_SCHEDULE	Implementation defined
<i>def-sched-var</i>	(none)	Implementation defined
<i>bind-var</i>	OMP_PROC_BIND	Implementation defined
<i>stacksize-var</i>	OMP_STACKSIZE	Implementation defined
<i>wait-policy-var</i>	OMP_WAIT_POLICY	Implementation defined
<i>thread-limit-var</i>	OMP_THREAD_LIMIT	Implementation defined
<i>max-active-levels-var</i>	OMP_MAX_ACTIVE_LEVELS	See description below
<i>active-levels-var</i>	(none)	<i>zero</i>
<i>levels-var</i>	(none)	<i>zero</i>
<i>place-partition-var</i>	OMP_PLACES	Implementation defined
<i>cancel-var</i>	OMP_CANCELLATION	<i>false</i>
<i>default-device-var</i>	OMP_DEFAULT_DEVICE	Implementation defined
<i>max-task-priority-var</i>	OMP_MAX_TASK_PRIORITY	<i>zero</i>
<i>tool-var</i>	OMP_TOOL	<i>enabled</i>
<i>tool-libraries-var</i>	OMP_TOOL_LIBRARIES	<i>empty string</i>
<i>def-allocator-var</i>	OMP_ALLOCATOR	Implementation defined

1

2

Description

3

- Each device has its own ICVs.

4

- The value of the *nthreads-var* ICV is a list.

5

- The value of the *bind-var* ICV is a list.

6

- The initial value of *dyn-var* is implementation defined if the implementation supports dynamic adjustment of the number of threads; otherwise, the initial value is *false*.

7

TABLE 2.2: Ways to Modify and to Retrieve ICV Values

ICV	Ways to modify value	Ways to retrieve value
<i>dyn-var</i>	<code>omp_set_dynamic()</code>	<code>omp_get_dynamic()</code>
<i>nest-var</i>	<code>omp_set_nested()</code>	<code>omp_get_nested()</code>
<i>nthreads-var</i>	<code>omp_set_num_threads()</code>	<code>omp_get_max_threads()</code>
<i>run-sched-var</i>	<code>omp_set_schedule()</code>	<code>omp_get_schedule()</code>
<i>def-sched-var</i>	(none)	(none)
<i>bind-var</i>	(none)	<code>omp_get_proc_bind()</code>
<i>stacksize-var</i>	(none)	(none)
<i>wait-policy-var</i>	(none)	(none)
<i>thread-limit-var</i>	<code>thread_limit</code> clause	<code>omp_get_thread_limit()</code>
1 <i>max-active-levels-var</i>	<code>omp_set_max_active_levels()</code>	<code>omp_get_max_active_levels()</code>
<i>active-levels-var</i>	(none)	<code>omp_get_active_level()</code>
<i>levels-var</i>	(none)	<code>omp_get_level()</code>
<i>place-partition-var</i>	(none)	See description below
<i>cancel-var</i>	(none)	<code>omp_get_cancellation()</code>
<i>default-device-var</i>	<code>omp_set_default_device()</code>	<code>omp_get_default_device()</code>
<i>max-task-priority-var</i>	(none)	<code>omp_get_max_task_priority()</code>
<i>tool-var</i>	(none)	(none)
<i>tool-libraries-var</i>	(none)	(none)
<i>def-allocator-var</i>	<code>omp_set_default_allocator()</code>	<code>omp_get_default_allocator()</code>

2

Description

- 3 • The value of the *nthreads-var* ICV is a list. The runtime call `omp_set_num_threads()` sets
4 the value of the first element of this list, and `omp_get_max_threads()` retrieves the value
5 of the first element of this list.
- 6 • The value of the *bind-var* ICV is a list. The runtime call `omp_get_proc_bind()` retrieves
7 the value of the first element of this list.
- 8 • Detailed values in the *place-partition-var* ICV are retrieved using the runtime calls
9 `omp_get_partition_num_places()`, `omp_get_partition_place_nums()`,
10 `omp_get_place_num_procs()`, and `omp_get_place_proc_ids()`.

1 **TABLE 2.3:** Scopes of ICVs

ICV	Scope
<i>dyn-var</i>	data environment
<i>nest-var</i>	data environment
<i>nthreads-var</i>	data environment
<i>run-sched-var</i>	data environment
<i>def-sched-var</i>	device
<i>bind-var</i>	data environment
<i>stacksize-var</i>	device
<i>wait-policy-var</i>	device
<i>thread-limit-var</i>	data environment
<i>max-active-levels-var</i>	device
<i>active-levels-var</i>	data environment
<i>levels-var</i>	data environment
<i>place-partition-var</i>	implicit task
<i>cancel-var</i>	global
<i>default-device-var</i>	data environment
<i>max-task-priority-var</i>	global
<i>tool-var</i>	global
<i>tool-libraries-var</i>	global
<i>def-allocator-var</i>	data environment

3 **Description**

- 4 • There is one copy per device of each ICV with device scope
- 5 • Each data environment has its own copies of ICVs with data environment scope
- 6 • Each implicit task has its own copy of ICVs with implicit task scope
- 7 Calls to OpenMP API routines retrieve or modify data environment scoped ICVs in the data
- 8 environment of their binding tasks.

table continued from previous page

ICV	construct clause, if used
<i>def-sched-var</i>	schedule
<i>bind-var</i>	proc_bind
<i>stacksize-var</i>	(none)
<i>wait-policy-var</i>	(none)
<i>thread-limit-var</i>	(none)
<i>max-active-levels-var</i>	(none)
1 <i>active-levels-var</i>	(none)
<i>levels-var</i>	(none)
<i>place-partition-var</i>	(none)
<i>cancel-var</i>	(none)
<i>default-device-var</i>	(none)
<i>max-task-priority-var</i>	(none)
<i>tool-var</i>	(none)
<i>tool-libraries-var</i>	(none)
<i>def-allocator-var</i>	(none)

2

Description

3

- The **num_threads** clause overrides the value of the first element of the *nthreads-var* ICV.

4

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

5

6

Cross References

7

- **parallel** construct, see Section 2.6 on page 54.

8

- **proc_bind** clause, Section 2.6 on page 54.

9

- **num_threads** clause, see Section 2.6.1 on page 59.

10

- Loop construct, see Section 2.8.1 on page 66.

11

- **schedule** clause, see Section 2.8.1.1 on page 74.

3.2.2 Memory spaces and allocators

1 **2.5 Memory Spaces and Allocators**2 **2.5.1 Memory Spaces**

3 OpenMP memory spaces represent storage where variables are defined. A set of
 4 memory traits and the value that those traits have define the characteristics of each
 5 memory space. Table 2.5 shows the supported memory traits, the possible values each
 6 trait can take and their meaning. Trait values and their names are not case sensitive.

TABLE 2.5: Memory traits and their values

Memory trait	Matching rule	Allowed values	Description
distance	≈	near, far	Specifies the relative physical distance of the memory space with respect to the task the request binds to.
bandwidth	≈	highest, lowest	Specifies the relative bandwidth of the memory space with respect to other memories in the system.
latency	≈	highest, lowest	Specifies the relative latency of the memory space with respect to other memories in the system.
location	=	see Table 2.6	Specifies the physical location of the memory space.

table continued on next page

table continued from previous page

Memory trait	Matching rule	Allowed values	Description
optimized	=	bandwidth, latency, capacity, none	Specifies if the memory space underlying technology is optimized to maximize a certain characteristic. The exact mapping of these values to actual technologies is implementation defined.
pagesize	=	positive integer	Specifies the size of the pages used by the memory space.
permission	=	r, w, rw	Specifies if read operations (r), write operations (w) or both (rw) are supported by the memory space.
capacity	≥	positive integer	Specifies the physical capacity in bytes of the memory space.
available	≥	positive integer	Specifies the current available capacity for new allocations in the memory space.

1

2

3

4

Table 2.6 shows the possible values for the **location** memory trait and their description. The values are not case sensitive. In addition, the **location** memory trait may accept other implementation specific values.

TABLE 2.6: Allowed values for the **location** memory trait

Location	Description
core	The memory space corresponds to a memory that is located within a core and might only be accessible by the hardware threads of that core.
socket	The memory space corresponds to a memory that is located within a socket and might only be accessible by the hardware threads of that socket.
device	The memory space corresponds to a memory that is located within the device and is accessible by any hardware thread of that device.

Certain constructs and API routines will try to find a memory space that matches a list of pairs of memory traits and values. A memory space matches a list if every trait in the list matches the corresponding trait in the memory spaces according to the following rules:

- An empty list of memory traits matches any memory space.
- Traits with the \geq matching rule match if the value of the trait in the memory space is greater or equal than the value in the list.
- Traits with the $=$ matching rule match if the value of the trait in the memory space is the same as the one in the list. For the `location` trait, for the matching to succeed it requires in addition that the task that the matching process binds to can access the memory space.
- Traits with the \approx matching rule match if the value of the trait in the memory space compared to the value of the trait in other candidate memory spaces results in the value in the list.
- The matching process selects first memory spaces that match the \geq and $=$ rules. From those selected in the previous step, it will select those that match the \approx rules.

If more than one memory space would match a memory trait specification it is unspecified which memory space will be returned by the matching process. If a list contains more than a pair with the same memory trait it is unspecified which memory space, if any, will be matched.

2.5.2 How Allocation Works

1 Allocations are made through requests to an allocator. Allocators can be either explicit,
 2 those created with the API calls defined in Section 3.5, or implicit, those logically created
 3 because of a construct. When an allocator receives a request to allocate storage of a
 4 certain size, it will try to return an allocation of logically consecutive virtual memory in its
 5 associated memory space of at least the size being requested. The behavior of the
 6 allocation process can be affected by the allocator traits that the user specifies. Table 2.7
 7 shows the allowed allocator traits, their possible values and the default value of each trait.
 8 Trait names and their values are not case sensitive.

TABLE 2.7: Allocator traits and their values

Allocator trait	Allowed values	Default value
<code>threadmodel</code>	<code>shared</code> , <code>exclusive</code>	<code>shared</code>
<code>alignment</code>	0, power of two integer	0
<code>pinned</code>	<code>true</code> , <code>false</code>	<code>false</code>
<code>fallback</code>	<code>null_fb</code> , <code>abort_fb</code> , <code>allocator_fb</code> , <code>default_fb</code>	<code>default_fb</code>
<code>fb_data</code>	an allocator handle	-

9
 10 When an allocator `threadmodel` trait is defined to be `exclusive` the implementation
 11 can assume that no operation will be performed on the allocator by more than one thread
 12 at a time.

13 If either the allocator `alignment` trait or the allocation alignment of the request is greater
 14 than zero the allocated memory will be byte aligned to the maximum of the two values.

15 When an allocator `pinned` trait is defined to be `true` then the allocated memory must be
 16 pinned to physical pages. If the `pinned` trait is defined to be `false` then the allocated
 17 memory needs not to be pinned to physical pages.

18 The `fallback` trait specifies how the allocator behaves when it cannot fulfil the allocation
 19 request. If the `fallback` trait is set to `null_fb` the allocator returns the value zero if fails
 20 to allocate the memory. If the `fallback` trait is set to `abort_fb` the program execution
 21 will be terminated if the allocation fails. If the `fallback` trait is set to `allocator_fb`
 22 then when an allocation fails the request will be delegated to the allocator specified in the
 23 `fb_data` trait. If the `fallback` trait is set to `default_fb` then when an allocation fails
 24 another allocation will be tried in a memory space with the
 25 `omp_default_memspace_traits` memory traits assuming all allocator traits to be set
 26 to their default values except for `fallback` which will be set to `null_fb`.

3.2.3 Changes to existing directives

▼ Fortran ▼

1 If any operation of the base language causes a reallocation of an array that is allocated
 2 with an explicit or implicit OpenMP allocator then that allocator will be used to release the
 3 current memory and to allocate the new memory.

▲ Fortran ▲

4 2.6 parallel Construct

5 Summary

6 This fundamental construct starts parallel execution. See Section 1.3 on page 15 for a general
 7 description of the OpenMP execution model.

8 Syntax

▼ C / C++ ▼

9 The syntax of the **parallel** construct is as follows:

```
#pragma omp parallel [clause[ [, ] clause] ... ] new-line
    structured-block
```

10 where *clause* is one of the following:

```
11     if ([parallel :] scalar-expression)
12     num_threads (integer-expression)
13     default (shared | none)
14     private (list)
15     firstprivate (list)
16     shared (list)
17     copyin (list)
18     reduction (reduction-identifier : list)
19     proc_bind (master | close | spread)
20     allocate (modifiers: list)
```

21



1 The syntax of the **parallel** construct is as follows:

```
!$omp parallel [clause[ [, ] clause]... ]
    structured-block
!$omp end parallel
```

2 where *clause* is one of the following:

```
3     if([parallel :] scalar-logical-expression)
4     num_threads (scalar-integer-expression)
5     default (private | firstprivate | shared | none)
6     private (list)
7     firstprivate (list)
8     shared (list)
9     copyin (list)
10    reduction (reduction-identifier : list)
11    proc_bind (master | close | spread)
12    allocate (modifiers: list)
```

13

14 The **end parallel** directive denotes the end of the **parallel** construct.



15 Binding

16 The binding thread set for a **parallel** region is the encountering thread. The encountering thread
17 becomes the master thread of the new team.

1 2.8.1 Loop Construct

2 Summary

3 The loop construct specifies that the iterations of one or more associated loops will be executed in
 4 parallel by threads in the team in the context of their implicit tasks. The iterations are distributed
 5 across threads that already exist in the team executing the **parallel** region to which the loop
 6 region binds.

7 Syntax

C / C++

8 The syntax of the loop construct is as follows:

```
#pragma omp for [clause[ [, ] clause]... ] new-line
for-loops
```

9 where clause is one of the following:

```
10 private (list)
11 firstprivate (list)
12 lastprivate ([ lastprivate-modifier : ] list)
13 linear (list[ : linear-step])
14 reduction (reduction-identifier : list)
15 schedule ([modifier [, modifier]:]kind[ , chunk_size])
16 collapse (n)
17 ordered[ (n) ]
18 nowait
19 allocate(modifiers:list)
```

21 The **for** directive places restrictions on the structure of all associated *for-loops*. Specifically, all
 22 associated *for-loops* must have *canonical loop form* (see Section 2.7 on page 62).

C / C++

Fortran

1 The syntax of the loop construct is as follows:

```
!$omp do [clause[ [, ] clause] ... ]
      do-loops
[$omp end do [nowait]]
```

2 where *clause* is one of the following:

```
3     private (list)
4     firstprivate (list)
5     lastprivate ([ lastprivate-modifier : ] list)
6     linear (list[ : linear-step])
7     reduction (reduction-identifier : list)
8     schedule ([modifier [, modifier] : ]kind[ , chunk_size])
9     collapse (n)
10    ordered[ (n) ]
11    allocate(modifiers:list)
```

13 If an **end do** directive is not specified, an **end do** directive is assumed at the end of the *do-loops*.

14 Any associated *do-loop* must be a *do-construct* or an *inner-shared-do-construct* as defined by the
 15 Fortran standard. If an **end do** directive follows a *do-construct* in which several loop statements
 16 share a **DO** termination statement, then the directive can only be specified for the outermost of these
 17 **DO** statements.

18 If any of the loop iteration variables would otherwise be shared, they are implicitly made private on
 19 the loop construct.

Fortran

20 Binding

21 The binding thread set for a loop region is the current team. A loop region binds to the innermost
 22 enclosing **parallel** region. Only the threads of the team executing the binding **parallel**
 23 region participate in the execution of the loop iterations and the implied barrier of the loop region if
 24 the barrier is not eliminated by a **nowait** clause.

1 2.8.2 sections Construct

2 Summary

3 The **sections** construct is a non-iterative worksharing construct that contains a set of structured
4 blocks that are to be distributed among and executed by the threads in a team. Each structured
5 block is executed once by one of the threads in the team in the context of its implicit task.

6 Syntax

▼ C / C++ ▼

7 The syntax of the **sections** construct is as follows:

```
#pragma omp sections [clause [ , ] clause] ... ] new-line
{
  [#pragma omp section new-line]
  structured-block
  [#pragma omp section new-line]
  structured-block]
  ...
}
```

8 where *clause* is one of the following:

```
9     private (list)
10    firstprivate (list)
11    lastprivate ([ lastprivate-modifier : ] list)
12    reduction (reduction-identifier : list)
13    nowait
14    allocate(modifiers:list)
```

15

▲ C / C++ ▲

 Fortran

1 The syntax of the **sections** construct is as follows:

```

!$omp sections [clause [ , ] clause] ... ]
  [!$omp section]
    structured-block
  [!$omp section
    structured-block]
  ...
!$omp end sections [nowait]
  
```

2 where *clause* is one of the following:

```

3     private (list)
4     firstprivate (list)
5     lastprivate ([ lastprivate-modifier : ] list)
6     reduction (reduction-identifier : list)
7     allocate(modifiers:list)
  
```

 Fortran

9 Binding

10 The binding thread set for a **sections** region is the current team. A **sections** region binds to
 11 the innermost enclosing **parallel** region. Only the threads of the team executing the binding
 12 **parallel** region participate in the execution of the structured blocks and the implied barrier of
 13 the **sections** region if the barrier is not eliminated by a **nowait** clause.

14 Description

15 Each structured block in the **sections** construct is preceded by a **section** directive except
 16 possibly the first block, for which a preceding **section** directive is optional.

17 The method of scheduling the structured blocks among the threads in the team is implementation
 18 defined.

19 There is an implicit barrier at the end of a **sections** construct unless a **nowait** clause is
 20 specified.

1 2.8.3 **single** Construct

2 **Summary**

3 The **single** construct specifies that the associated structured block is executed by only one of the
 4 threads in the team (not necessarily the master thread), in the context of its implicit task. The other
 5 threads in the team, which do not execute the block, wait at an implicit barrier at the end of the
 6 **single** construct unless a **nowait** clause is specified.

7 **Syntax**

8 The syntax of the **single** construct is as follows: C / C++

```
#pragma omp single [clause[ [, ] clause] ... ] new-line
    structured-block
```

9 where *clause* is one of the following:

10 **private** (*list*)
 11 **firstprivate** (*list*)
 12 **copyprivate** (*list*)
 13 **nowait**
 14 **allocate**(*modifiers:list*)

15 C / C++
Fortran

16 The syntax of the **single** construct is as follows:

```
!$omp single [clause[ [, ] clause] ... ]
    structured-block
!$omp end single [end_clause[ [, ] end_clause] ... ]
```

17 where *clause* is one of the following:

```

1      private(list)
2      firstprivate(list)
3      allocate(modifiers:list)

```

```
4
```

5 and *end_clause* is one of the following:

```
6      copyprivate(list)
```

```
7      nowait
```

▲────────────────────────────────── Fortran ───────────────────────────────────▲

8 Binding

9 The binding thread set for a **single** region is the current team. A **single** region binds to the
10 innermost enclosing **parallel** region. Only the threads of the team executing the binding
11 **parallel** region participate in the execution of the structured block and the implied barrier of the
12 **single** region if the barrier is not eliminated by a **nowait** clause.

13 Description

14 The method of choosing a thread to execute the structured block is implementation defined. There
15 is an implicit barrier at the end of the **single** construct unless a **nowait** clause is specified.

16 Events

17 The *single-begin* event occurs after an **implicit task** encounters a **single** construct but
18 before the task starts the execution of the structured block of the **single** region.

19 The *single-end* event occurs after a **single** region finishes execution of the structured block but
20 before resuming execution of the encountering implicit task.

21 Tool Callbacks

22 A thread dispatches a registered **ompt_callback_work** callback for each occurrence of
23 *single-begin* and *single-end* events in that thread. The callback has type signature
24 **ompt_callback_work_t**. The callback receives **ompt_scope_begin** or
25 **ompt_scope_end** as its *endpoint* argument, as appropriate, and
26 **ompt_work_single_executor** or **ompt_work_single_other** as its *wstype* argument.

1 2.10 Tasking Constructs

2 2.10.1 task Construct

3 Summary

4 The **task** construct defines an explicit task.

5 Syntax

C / C++

6 The syntax of the **task** construct is as follows:

```
#pragma omp task [clause[ [, ] clause] ... ] new-line
    structured-block
```

7 where *clause* is one of the following:

```
8     if([ task :] scalar-expression)
9     final (scalar-expression)
10    untied
11    default (shared | none)
12    mergeable
13    private (list)
14    firstprivate (list)
15    shared (list)
16    in_reduction (reduction-identifier : list)
17    depend (dependence-type : locator-list)
18    priority (priority-value)
19    allocate(modifiers:list)
```

20

C / C++


 Fortran

1 The syntax of the **task** construct is as follows:

```

!$omp task [clause [ , ] clause ... ]
           structured-block
!$omp end task
  
```

2 where *clause* is one of the following:

```

3     if([ task : ] scalar-logical-expression)
4     final (scalar-logical-expression)
5     untied
6     default(private | firstprivate | shared | none)
7     mergeable
8     private (list)
9     firstprivate (list)
10    shared (list)
11    in_reduction (reduction-identifier : list)
12    depend (dependence-type : locator-list)
13    priority (priority-value)
14    allocate(modifiers:list)
  
```


 Fortran

16 Binding

17 The binding thread set of the **task** region is the current team. A **task** region binds to the
 18 innermost enclosing **parallel** region.

```

1      num_tasks (num-tasks)
2      collapse (n)
3      final (scalar-expr)
4      priority (priority-value)
5      untied
6      mergeable
7      nogroup
8      allocate(modifiers:list)

```

10 The **taskloop** directive places restrictions on the structure of all associated *for-loops*.
 11 Specifically, all associated *for-loops* must have canonical loop form (see Section 2.7 on page 62).



12 The syntax of the **taskloop** construct is as follows:

```

!$omp taskloop [clause[[,] clause] ...]
    do-loops
[!$omp end taskloop]

```

13 where *clause* is one of the following:

```

14      if([ taskloop :] scalar-logical-expr)
15      shared(list)
16      private(list)
17      firstprivate(list)
18      lastprivate(list)
19      reduction(reduction-identifier : list)
20      in_reduction(reduction-identifier : list)
21      default(private | firstprivate | shared | none)
22      grainsize(grain-size)
23      num_tasks(num-tasks)
24      collapse(n)

```

```

1      final (scalar-logical-expr)
2      priority (priority-value)
3      untied
4      mergeable
5      nogroup
6      allocate(modifiers:list)

```

7

8 If an **end taskloop** directive is not specified, an **end taskloop** directive is assumed at the end
9 of the *do-loops*.

10 Any associated *do-loop* must be *do-construct* or an *inner-shared-do-construct* as defined by the
11 Fortran standard. If an **end taskloop** directive follows a *do-construct* in which several loop
12 statements share a **DO** termination statement, then the directive can only be specified for the
13 outermost of these **DO** statements.

14 If any of the loop iteration variables would otherwise be shared, they are implicitly made private for
15 the loop-iteration tasks generated by the **taskloop** construct. Unless the loop iteration variables
16 are specified in a **lastprivate** clause on the **taskloop** construct, their values after the loop
17 are unspecified.

▲────────────────────────────────── Fortran ───────────────────────────────────▲

18 Binding

19 The binding thread set of the **taskloop** region is the current team. A **taskloop** region binds to
20 the innermost enclosing **parallel** region.

21 Description

22 The **taskloop** construct is a *task generating construct*. When a thread encounters a **taskloop**
23 construct, the construct partitions the associated loops into explicit tasks for parallel execution of
24 the loops' iterations. The data environment of each generated task is created according to the
25 data-sharing attribute clauses on the **taskloop** construct, per-data environment ICVs, and any
26 defaults that apply. The order of the creation of the loop tasks is unspecified. Programs that rely on
27 any execution order of the logical loop iterations are non-conforming.

28 By default, the **taskloop** construct executes as if it was enclosed in a **taskgroup** construct
29 with no statements or directives outside of the **taskloop** construct. Thus, the **taskloop**
30 construct creates an implicit **taskgroup** region. If the **nogroup** clause is present, no implicit
31 **taskgroup** region is created.

3.2.4 Allocate directive and clause

1 As another example, if a lock acquire and release happen in different parts of a task region, no
 2 attempt should be made to acquire the same lock in any part of another task that the executing
 3 thread may schedule. Otherwise, a deadlock is possible. A similar situation can occur when a
 4 **critical** region spans multiple parts of a task and another schedulable task contains a
 5 **critical** region with the same name.

6 The use of threadprivate variables and the use of locks or critical sections in an explicit task with an
 7 **if** clause must take into account that when the **if** clause evaluates to *false*, the task is executed
 8 immediately, without regard to *Task Scheduling Constraint 2*.



9 Events

10 The *task-schedule* event occurs in a thread when the thread switches tasks at a task scheduling
 11 point; no event occurs when switching to or from a merged task.

12 Tool Callbacks

13 A thread dispatches a registered **ompt_callback_task_schedule** callback for each
 14 occurrence of a *task-schedule* event in the context of the task that begins or resumes. This callback
 15 has the type signature **ompt_callback_task_schedule_t**. The argument *prior_task_status*
 16 is used to indicate the cause for suspending the prior task. This cause may be the completion of the
 17 prior task region, the encountering of a **taskyield** construct, or the encountering of an active
 18 cancellation point.

19 Cross References

20 • **ompt_callback_task_schedule_t**, see Section [4.6.2.10](#) on page [409](#).

21 2.11 Memory Management Directives

22 2.11.1 allocate Directive

23 Summary

24 The **allocate** directive specifies how a set of variables are allocated.



25 The **allocate** directive is a declarative directive.

▲────────────────────────────────── C / C++ ───────────────────────────────────▲
 ▼────────────────────────────────── Fortran ───────────────────────────────────▼

1 The **allocate** directive is a declarative directive if it is not associated with an **allocate**
 2 statement.

▲────────────────────────────────── Fortran ───────────────────────────────────▲

3 **Syntax**

▼────────────────────────────────── C / C++ ───────────────────────────────────▼

4 The syntax of the **allocate** directive is as follows:

```
#pragma omp allocate(list) [clause[ [ [, ] clause] ... ]] new-line
```

5 where clause is one of the following:

6 allocator(allocator)

7 memspace(memspace)

8 alloctrails(alloctrail-list)

9 memtraits(memtrait-list)

10 safe_align(alignment)

11 where allocator is an expression of the **omp_allocator_t** type.

12 where memspace is an expression of the **omp_memspace_t** type.

13 where alloctrail-list is a key-value list where the allowed keys are the allocator traits keys
 14 and the allowed values are the accepted values for each key.

15 where memtrait-list is a key-value list where the allowed keys are the memory traits keys
 16 and the allowed values are the accepted values for each key.

17 where alignment is an integer expression that must evaluate to a power of two.

▲────────────────────────────────── C / C++ ───────────────────────────────────▲

▼────────────────────────────────── Fortran ───────────────────────────────────▼

18 The syntax of the **allocate** directive is as follows:

```
!$omp allocate(list) [clause[ [ [, ] clause] ... ]]
```

19 or

```
!$omp allocate [ (list) ] clause [ [ [, ] clause ] ... ]
allocate statement
```

1 where clause is one of the following:

- 2 allocator(allocator)
- 3 memspace(memspace)
- 4 alloctrails(alloctrail-list)
- 5 memtraits(memtrait-list)
- 6 safe_align(alignment)

7 where allocator is an integer expression of the `omp_allocator_kind` kind.

8 where memspace is an integer expression of the `omp_memspace_kind` kind.

9 where alloctrail-list is a key-value list where the allowed keys are the allocator traits keys
10 and the allowed values are the accepted values for each key.

11 where memtrait-list is a key-value list where the allowed keys are the memory traits keys
12 and the allowed values are the accepted values for each key.

13 where alignment is an integer expression that must evaluate to a power of two.

▲────────────────── Fortran ───────────────────▲

14 **Description**

15 If the directive is not associated with a Fortran `allocate` statement, the storage for each
16 list item that appears in the directive will be provided by an allocation through an allocator.
17 If no clause is specified then the allocator specified by the `def-allocator-var` ICV will be
18 used. If the `allocator` clause is specified, the allocator specified in the clause will be used.
19 Otherwise, the allocation will be provided as if using an allocator that had been built with
20 the specified allocator traits, memory traits and/or the memspace memory space. If the
21 `safe_align` clause is specified, then the allocation alignment of the request will be the
22 value of the `safe_align` clause.

23 The scope of this allocation is that of the list item in the base language. When the
24 allocation reaches the end of the scope it will be deallocated through the specified
25 allocator or as if using an allocator that had been built with the specified allocator traits,
26 memory traits and/or the memspace memory space. If the execution leaves the scope in a
27 manner not supported by the base language it is unspecified whether the deallocation
28 happens or not.

▼────────────────── Fortran ───────────────────▶

1 If the directive is associated with a Fortran `allocate` statement, the allocation of the
 2 specified list items will be provided through an allocator. If no clause is specified then the
 3 allocator specified by the `def-allocator-var` ICV will be used. If the `allocator` clause is
 4 specified, the allocator specified in the clause will be used. Otherwise, the allocation will
 5 be provided as if using an allocator that had been built with the specified `allocator` traits,
 6 memory traits and/or the `memspace` memory space. If no list item is specified then all
 7 variables allocated by the `allocate` statement will be provided by the allocator.

▲────────────────── Fortran ───────────────────▶

8 For allocations that arise from this directive the `null_fb` value of the `fallback` allocator
 9 trait will behave as if the `abort_fb` had been specified.

Restrictions

- 11 • A variable that is part of another variable (as an array or structure element) cannot
 12 appear in an `allocate` directive.
- 13 • The directive must appear in the same scope of the list item declaration and before its
 14 first use.
- 15 • If the `allocator` clause is present, no other clause must be specified.
- 16 • If the `allocator` clause is present, the allocator must be an allocator returned by the
 17 `omp_init_allocator` routine.
- 18 • At most one `allocator` clause can appear on the `allocate` directive.
- 19 • If the `memspace` clause is present, the `memtraits` clause must not be specified.
- 20 • If the `memspace` clause is present, the memspace must be a memory space returned
 21 by the `omp_init_memspace` routine.
- 22 • At most one `memspace` clause can appear on the `allocate` directive.
- 23 • If the `safe_align` clause is present, its value must a power of two.

▼────────────────── C / C++ ───────────────────▶

- 24 • If a list item has a static storage type, the `allocator` and the `memspace` clauses must
 25 not be specified.
- 26 • If a list item has a static storage type, the `fallback` allocator trait must not have the
 27 `allocator_fb` value.

▲────────────────── C / C++ ───────────────────▶

▼ Fortran ▼

- 1 • List items specified in the `allocate` directive must not have the `ALLOCATABLE`
- 2 attribute unless the directive is associated with an `allocate` statement.
- 3 • List items specified in an `allocate` directive that is associated with an `allocate`
- 4 statement must be `ALLOCATABLE` variables allocated by the `allocate` statement.

▲ Fortran ▲

5 Cross References

- 6 • Memory spaces, allocators and their traits, see Section 2.5 on page 50.

▼ C / C++ ▼

- 7 • `omp_memspace_t` and `omp_allocator_t`, see Section 3.5.1 on page 328.

▲ C / C++ ▲

▼ Fortran ▼

- 8 • `omp_memspace_kind` and `omp_allocator_kind`, see Section 3.5.1 on page 328.

▲ Fortran ▲

9 2.11.2 The `allocate` Clause

10 Summary

11 The `allocate` clause specifies the allocation and memory traits of the storage used for
 12 private variables of a directive.

1 **Syntax**2 The syntax of the `allocate` clause is as follows:

allocate (*[modifiers:] list*)

3 where `modifiers` is a comma separated list of one or more of the following:4 `allocator(allocator)`5 `memspace(memspace)`6 `alloctrails(alloctrail-list)`7 `memtraits(memtrait-list)`8 `safe_align(alignment)`

C / C++

9 where `allocator` is an integer expression of the `omp_allocator_t` type.10 where `memspace` is an integer expression of the `omp_memspace_t` type.

C / C++

Fortran

11 where `allocator` is an integer expression of the `omp_allocator_kind` kind.12 where `memspace` is an integer expression of the `omp_memspace_kind` kind.

Fortran

13 where `alloctrail-list` is a key-value list where the allowed keys are the allocator traits keys
14 and the allowed values are the accepted values for each key.15 where `memtrait-list` is a key-value list where the allowed keys are the memory traits keys
16 and the allowed values are the accepted values for each key.17 where `alignment` is an integer expression that must evaluate to a power of two.

Description

The storage for new list items that arise from list item that appear in the directive will be provided by an allocation through an allocator. If no modifier is specified then the allocator specified by the `def-allocator-var` ICV will be used. If the allocator modifier is specified, the allocator specified in the clause will be used. Otherwise, the allocation will be provided as if using an allocator that had been built with the specified allocator traits, memory traits and/or the memspace memory space. For allocations that arise from this clause the `null_fb` value of the `fallback` allocator trait will behave as if the `abort_fb` had been specified. If the `safe_align` modifier is specified, then the allocation alignment of the request will be the value of the `safe_align` modifier.

Restrictions

- List items specified in the `allocate` clause must also be specified in a `private`, `firstprivate`, `lastprivate`, `linear` or `reduction` clause in the same directive.
- If the `allocator` modifier is present, no other modifier must be specified.
- If the `allocator` modifier is present, the allocator must be an allocator returned by the `omp_init_allocator` routine.
- At most one `allocator` modifier can appear on the `allocate` clause.
- If the `memspace` modifier is present, the `memtraits` modifier must not be specified.
- If the `memspace` modifier is present, the memspace must be a memory space returned by the `omp_init_memspace` routine.
- At most one `memspace` modifier can appear on the `allocate` modifier.

Cross References

- Memory spaces, allocators and their traits, see Section 2.5 on page 50.

▼ C / C++ ▼

- `omp_memspace_t` and `omp_allocator_t`, see Section 3.5.1 on page 328.

▲ C / C++ ▲

▼ Fortran ▼

- `omp_memspace_kind` and `omp_allocator_kind`, see Section 3.5.1 on page 328.

▲ Fortran ▲

3.3 Changes to Chapter 3

C / C++

1 3.5 Memory Management Routines

2 This section describes routines that support management of memory on the current
3 device.

4 Instances of OpenMP memory management types must be accessed only through the
5 routines described in this section; programs that otherwise access OpenMP instances of
6 these types are non-conforming.

7 3.5.1 Memory Management Types

8 The following type definitions are used by the memory management routines:

C / C++

9 The type `omp_uintptr_t` must be defined as an unsigned integer that is capable of
10 storing an address.

```
11 typedef enum {
12     OMP_MTK_DISTANCE,
13     OMP_MTK_LOCATION,
14     OMP_MTK_BANDWIDTH,
15     OMP_MTK_LATENCY,
16     OMP_MTK_OPTIMIZED
17     OMP_MTK_PAGESIZE,
18     OMP_MTK_PERMISSION,
19     OMP_MTK_CAPACITY,
20     OMP_MTK_AVAILABLE
21 } omp_memtrait_key_t;
22
23 typedef enum {
24     OMP_MTV_FALSE = 0,
25     OMP_MTV_TRUE = 1,
26     OMP_MTV_NEAR,
27     OMP_MTV_FAR,
28     OMP_MTV_CORE,
29     OMP_MTV_SOCKET,
30     OMP_MTV_NODE,
31     OMP_MTV_HIGHEST,
```



```
1         OMP_MTV_LOWEST,  
2         OMP_MTV_BANDWIDTH,  
3         OMP_MTV_LATENCY,  
4         OMP_MTV_CAPACITY,  
5         OMP_MTV_NONE,  
6         OMP_MTV_R,  
7         OMP_MTV_W,  
8         OMP_MTV_RW = OMP_MTV_R | OMP_MTV_W,  
9     } omp_memtrait_value_t;  
10  
11  
12     typedef struct {  
13         omp_memtrait_key_t key;  
14         omp_uintptr_t value;  
15     } omp_memtrait_t;  
16  
17  
18     typedef enum {  
19         OMP_ATK_THREADMODEL,  
20         OMP_ATK_ALIGNMENT,  
21         OMP_ATK_PIN,  
22         OMP_ATK_FALLBACK,  
23         OMP_ATK_FB_DATA  
24     } omp_alloctrain_key_t;  
25  
26  
27     typedef enum {  
28         OMP_ATV_FALSE = 0,  
29         OMP_ATV_TRUE = 1,  
30         OMP_ATV_SHARED,  
31         OMP_ATV_EXCLUSIVE,  
32         OMP_ATV_ABORT_FB,  
33         OMP_ATV_NULL_FB,  
34         OMP_ATV_ALLOCATOR_FB,  
35         OMP_ATV_DEFAULT_FB  
36     } omp_alloctrain_value_t;  
37  
38  
39     typedef struct {  
40         omp_alloctrain_key_t key;  
41         omp_uintptr_t value;  
42     } omp_alloctrain_t;  
43
```

```

1
2     omp_memtrait_set_t;
3     const omp_memtrait_set_t omp_default_memspace_traits;
4     omp_memspace_t;
5     enum { OMP_NULL_MEMSPACE = NULL };
6
7
8     omp_alloctrail_set_t;
9     const omp_alloctrail_set_t omp_default_allocator_traits;
10    omp_allocator_t;
11    enum { OMP_NULL_ALLOCATOR = NULL };
12
13
14    ▲────────────────────────────────── C / C++ ───────────────────────────────────►
15    ▼────────────────────────────────── Fortran ───────────────────────────────────►
16
17    integer parameter omp_memtrait_key_kind
18
19    integer(kind=omp_memtrait_key_kind), &
20        parameter :: omp_mtk_distance
21    integer(kind=omp_memtrait_key_kind), &
22        parameter :: omp_mtk_location
23    integer(kind=omp_memtrait_key_kind), &
24        parameter :: omp_mtk_bandwidth
25    integer(kind=omp_memtrait_key_kind), &
26        parameter :: omp_mtk_latency
27    integer(kind=omp_memtrait_key_kind), &
28        parameter :: omp_mtk_optimized
29    integer(kind=omp_memtrait_key_kind), &
30        parameter :: omp_mtk_pagesize
31    integer(kind=omp_memtrait_key_kind), &
32        parameter :: omp_mtk_permission
33    integer(kind=omp_memtrait_key_kind), &
34        parameter :: omp_mtk_capacity
35    integer(kind=omp_memtrait_key_kind), &
36        parameter :: omp_mtk_available
37
38    integer parameter omp_memtrait_val_kind
39
40    integer(kind=omp_memtrait_val_kind), &
41        parameter :: omp_mtv_false = 0
42    integer(kind=omp_memtrait_val_kind), &
43        parameter :: omp_mtv_true = 1

```

```
1      integer(kind=omp_memtrait_val_kind), &
2          parameter :: omp_mtv_near
3      integer(kind=omp_memtrait_val_kind), &
4          parameter :: omp_mtv_far
5      integer(kind=omp_memtrait_val_kind), &
6          parameter :: omp_mtv_core
7      integer(kind=omp_memtrait_val_kind), &
8          parameter :: omp_mtv_socket
9      integer(kind=omp_memtrait_val_kind), &
10         parameter :: omp_mtv_node
11     integer(kind=omp_memtrait_val_kind), &
12         parameter :: omp_mtv_highest
13     integer(kind=omp_memtrait_val_kind), &
14         parameter :: omp_mtv_lowest
15     integer(kind=omp_memtrait_val_kind), &
16         parameter :: omp_mtv_bandwidth
17     integer(kind=omp_memtrait_val_kind), &
18         parameter :: omp_mtv_latency
19     integer(kind=omp_memtrait_val_kind), &
20         parameter :: omp_mtv_capacity
21     integer(kind=omp_memtrait_val_kind), &
22         parameter :: omp_mtv_none
23     integer(kind=omp_memtrait_val_kind), &
24         parameter :: omp_mtv_r
25     integer(kind=omp_memtrait_val_kind), &
26         parameter :: omp_mtv_w
27     integer(kind=omp_memtrait_val_kind), &
28         parameter :: omp_mtv_rw = IOR(omp_mtv_r,omp_mtv_w)
29
30
31     type omp_memtrait
32         integer(kind=omp_memtrait_key_kind) key
33         integer(kind=omp_memtrait_val_kind) value
34     end type omp_memtrait
35
36
37     integer parameter omp_alloctrait_key_kind
38
39
40     integer(kind=omp_alloctrait_key_kind), &
41         parameter :: omp_atk_threadmodel
42     integer(kind=omp_alloctrait_key_kind), &
43         parameter :: omp_atk_alignment
```

```
1         integer(kind=omp_alloctratit_key_kind), &
2             parameter :: omp_atk_pin
3         integer(kind=omp_alloctratit_key_kind), &
4             parameter :: omp_atk_fallback
5         integer(kind=omp_alloctratit_key_kind), &
6             parameter :: omp_atk_fb_data
7
8
9         integer parameter omp_alloctratit_val_kind
10
11
12         integer(kind=omp_alloctratit_val_kind), &
13             parameter :: omp_atv_true = 0
14         integer(kind=omp_alloctratit_val_kind), &
15             parameter :: omp_atv_false = 1
16         integer(kind=omp_alloctratit_val_kind), &
17             parameter :: omp_atv_shared
18         integer(kind=omp_alloctratit_val_kind), &
19             parameter :: omp_atv_exclusive
20         integer(kind=omp_alloctratit_val_kind), &
21             parameter :: omp_atv_abort_fb
22         integer(kind=omp_alloctratit_val_kind), &
23             parameter :: omp_atv_null_fb
24         integer(kind=omp_alloctratit_val_kind), &
25             parameter :: omp_atv_allocator_fb
26         integer(kind=omp_alloctratit_val_kind), &
27             parameter :: omp_atv_default_fb
28
29
30         type omp_alloctratit
31             integer(kind=omp_alloctratit_key_kind) key
32             integer(kind=omp_alloctratit_val_kind) value
33         end type omp_alloctratit
34
35
36         integer parameter omp_memtrait_set_kind
37         integer(kind=omp_memtrait_set_kind), &
38             parameter :: omp_default_memspace_traits
39         integer parameter omp_memspace_kind
40         integer(kind=omp_memspace_kind), &
41             parameter :: omp_null_memspace = 0
42
43
```

3.3.1 Routines for defining memory traits

```

1      integer parameter omp_alloctrail_set_kind
2      integer(kind=omp_alloctrail_set_kind), &
3          parameter :: omp_default_allocator_traits
4      integer parameter omp_allocator_kind
5      integer(kind=omp_allocator_kind), &
6          parameter :: omp_null_allocator = 0

```

▲────────────────── Fortran ───────────────────▲

7 3.5.2 omp_init_memtrait_set

8 Summary

9 The `omp_init_memtrait_set` routine initializes an OpenMP memory traits set.

10 Format

▼────────────────── C / C++ ───────────────────▼

```

void omp_init_memtrait_set (omp_memtrait_set_t *set,
                           size_t ntraits,
                           omp_memtrait_t *traits);           (C)
void omp_init_memtrait_set (omp_memtrait_set_t *set,
                           size_t ntraits = 0,
                           omp_memtrait_t *traits = NULL);    (C++)

```

▲────────────────── C / C++ ───────────────────▲

▼────────────────── Fortran ───────────────────▼

```

subroutine omp_init_memtrait_set ( set, ntraits, traits )
integer(kind=omp_memtrait_set_kind),intent(out) :: set
integer,intent(in) :: ntraits
type(omp_memtrait),intent(in) :: traits(*)

```

▲────────────────── Fortran ───────────────────▲

11 Binding

12 The binding thread set for an `omp_init_memtrait_set` region is all threads on a
13 device. The effect of executing this routine is not related to any specific region
14 corresponding to any construct or API routine.

1 **Constraints on Arguments**

2 If the `ntraits` argument is greater than zero, then there must be at least as many traits
3 specified in the `traits` argument; otherwise, the behavior is unspecified.

4 **Effect**

5 The effect of the `omp_init_memtrait_set` routine is to initialize the memory trait set in
6 the `set` argument to the memory traits specified in the `traits` argument. The number of
7 traits to be included in the set is specified by the `ntraits` argument.

8 **3.5.3 omp_destroy_memtrait_set**

9 **Summary**

10 The `omp_destroy_memtrait_set` routine ensures that an OpenMP memory traits set
11 is uninitialized.

12 **Format**

▼ C / C++ ▶

```
void omp_destroy_memtrait_set (omp_memtrait_set_t *set);
```

▲ C / C++ ▶

▼ Fortran ▶

```
subroutine omp_destroy_memtrait_set ( set )
integer(kind=omp_memtrait_set_kind),intent(inout) :: set
```

▲ Fortran ▶

13 **Binding**

14 The binding thread set for an `omp_destroy_memtrait_set` region is all threads on a
15 device. The effect of executing this routine is not related to any specific region
16 corresponding to any construct or API routine.

1 **Effect**

2 The effect of the `omp_destroy_memtrait_set` routine is to uninitialized the memory
3 traits set specified in the first argument.

4 **3.5.4 omp_add_memtraits**5 **Summary**

6 The `omp_add_memtraits` routine adds a memory trait to the memory traits set.

7 **Format**

▼────────────────── C / C++ ───────────────────▼

```
void omp_add_memtraits (omp_memtrait_set_t *set,
                       size_t ntraits,
                       omp_memtrait_t *traits);
```

▲────────────────── C / C++ ───────────────────▲

▼────────────────── Fortran ───────────────────▼

```
subroutine omp_add_memtraits ( set, ntraits, traits )
integer(kind=omp_memtrait_set_kind),intent(inout) :: set
integer,intent(in) :: ntraits
type(omp_memtrait),intent(in) :: traits(*)
```

▲────────────────── Fortran ───────────────────▲

8 **Constraints on Arguments**

9 If the `ntraits` argument is greater than zero, then there must be at least as many traits
10 specified in the `traits` argument; otherwise, the behavior is unspecified.

11 **Binding**

12 The binding thread set for an `omp_add_memtraits` region is all threads on a device. The
13 effect of executing this routine is not related to any specific region corresponding to any
14 construct or API routine.

1 **Effect**

2 The effect of the `omp_add_memtraits` routine is that the ntraits specified in traits are
3 added to the set of memory traits.

4 **Cross References**

- 5 • Memory traits in Section 2.5.1 on page 50

6 **3.5.5 omp_merge_memtraits**7 **Summary**

8 The `omp_merge_memtraits` routine merges two memory traits sets.

9 **Format**

▼ C / C++ ▼

```
void omp_merge_memtraits (omp_memtrait_set_t *dst,
                          const omp_memtrait_set_t *src,
                          int dst_priority ); (C)
void omp_merge_memtraits (omp_memtrait_set_t *dst,
                          const omp_memtrait_set_t *src,
                          bool dst_priority = true); (C++)
```

▲ C / C++ ▲

▼ Fortran ▼

```
subroutine omp_merge_memtraits ( dst, src, dst_priority )
integer(kind=omp_memtrait_set_kind),intent(inout) :: dst
integer(kind=omp_memtrait_set_kind),intent(in) :: src
logical :: dst_priority
```

▲ Fortran ▲

10 **Binding**

11 The binding thread set for an `omp_merge_memtraits` region is all threads on a device.
12 The effect of executing this routine is not related to any specific region corresponding to
13 any construct or API routine.

3.3.2 Routines for memory spaces

1 **Effect**

2 The effect of the `omp_merge_memtraits` routine is that the two memory traits sets `dst`
 3 and `src` are merged into `dst`. If the `available` trait appears in both sets the merged
 4 value for the trait will be the result of adding the values in each set. If the `capacity` trait
 5 appears in both sets the merged value for the trait will be the greater of the values in either
 6 set. For any other trait, if the same memory trait appears in both sets, if the `dst_priority`
 7 argument evaluates to `true` the merged value will be that of the `dst` set; otherwise, it will
 8 be the value of the `src` set.

9 **3.5.6 `omp_init_memspace`**

10 **Summary**

11 The `omp_init_memspace` routine returns handler to a memory space that matches the
 12 specified memory traits.

13 **Format**

▼ C / C++ ▼

```
omp_memspace_t * omp_init_memspace(const omp_memtrait_set_t *traits);
```

▲ C / C++ ▲

▼ Fortran ▼

```
integer(kind=omp_memspace_kind) function omp_init_memspace (traits)  
integer(kind=omp_memtrait_set_kind), intent(in) :: traits
```

▲ Fortran ▲

14 **Binding**

15 The binding thread set for an `omp_init_memspace` region is all threads on a device. The
 16 effect of executing this routine is not related to any specific region corresponding to any
 17 construct or API routine.

18 **Constraints on Arguments**

19 The `traits` argument must have been initialized with the `omp_init_memtrait_set`
 20 routine.

1 **Effect**

2 The `omp_init_memspace` routine returns a handler to a memory space in the current
 3 device that matches the memory traits specified in the traits set. If no memory space is
 4 found that matches the specified memory traits then the special value
 5 `OMP_NULL_MEMSPACE` is returned.

6 The traits in `omp_default_memspace_traits` must be defined in such a way that it
 7 guarantees that the `omp_init_memspace` routine will return a valid memory space that
 8 is always the same and that an allocation from that memory space is guaranteed to be
 9 accessible to all threads on that device without any special consideration.

10 **Cross References**

- 11
 - Memory spaces in Section 2.5.1 on page 50

12 **3.5.7 `omp_destroy_memspace`**13 **Summary**

14 The `omp_destroy_memspace` releases all resources associated with a memory space
 15 handler.

16 **Format**

▼ C / C++ ▼

```
void omp_destroy_memspace (omp_memspace_t *memspace);
```

▲ C / C++ ▲

▼ Fortran ▼

```
subroutine omp_destroy_memspace ( memspace )  
integer (kind=omp_memspace_kind), intent (out) :: memspace
```

▲ Fortran ▲

3.3.3 Routines for defining allocator traits

1 **Binding**

2 The binding thread set for an `omp_destroy_memspace` region is all threads on a device.
 3 The effect of executing this routine is not related to any specific region corresponding to
 4 any construct or API routine.

5 **Effect**

6 The `omp_destroy_memspace` routine releases resources associated with the
 7 memspace handler. Accessing allocators, or memory allocated by them, that have been
 8 associated through the memspace handler results in unspecified behavior.

9 **3.5.8 `omp_init_alloctrain_set`**

10 **Summary**

11 The `omp_init_alloctrain_set` initializes an OpenMP allocator traits set.

12 **Format**

▼ C / C++ ▲

```
void omp_init_alloctrain_set (omp_alloctrain_set_t *set
                             size_t ntraits,
                             omp_alloctrain_t *traits);      (C)
void omp_init_alloctrain_set (omp_alloctrain_set_t *set
                             size_t ntraits = 0,
                             omp_alloctrain_t *traits = NULL); (C++)
```

▲ C / C++ ▲

▼ Fortran ▲

```
subroutine omp_init_alloctrain_set ( set, ntraits, traits )
integer(kind=omp_alloctrain_set_kind),intent(out) :: set
integer,intent(in) :: ntraits
type(omp_alloctrain),intent(in) :: traits(*)
```

▲ Fortran ▲

1 **Binding**

2 The binding thread set for an `omp_init_alloctrail_set` region is all threads on a
 3 device. The effect of executing this routine is not related to any specific region
 4 corresponding to any construct or API routine.

5 **Constraints on Arguments**

6 If the `ntraits` argument is greater than zero, then there must be at least as many traits
 7 specified in the `traits` argument. If there are fewer than `ntraits` traits the behavior is
 8 unspecified.

9 **Effect**

10 The effect of the `omp_init_alloctrail_set` routine is to initialize the allocator trait set
 11 in the `set` argument to the allocator traits specified in the `traits` argument. The number of
 12 traits to be included in the set is specified by the `ntraits` argument.

13 **3.5.9 omp_destroy_alloctrail_set**

14 **Summary**

15 The `omp_destroy_alloctrail_set` routine ensures that an OpenMP allocator traits
 16 set is uninitialized.

17 **Format**

▼──────────────────────────────── C / C++ ─────────────────────────────────▶

```
void omp_destroy_alloctrail_set (omp_alloctrail_set_t *set);
```

▲──────────────────────────────── C / C++ ─────────────────────────────────▲

▼──────────────────────────────── Fortran ─────────────────────────────────▶

```
subroutine omp_destroy_alloctrail_set ( set )  
integer(kind=omp_alloctrail_set_kind),intent(inout) :: set
```

▲──────────────────────────────── Fortran ─────────────────────────────────▲

1 **Binding**

2 The binding thread set for an `omp_destroy_allocator_set` region is all threads on a
 3 device. The effect of executing this routine is not related to any specific region
 4 corresponding to any construct or API routine.

5 **Effect**

6 The effect of the `omp_destroy_allocator_set` routine is to uninitialized the allocator
 7 traits set specified in the first argument.

8 **3.5.10 omp_add_allocator_traits**9 **Summary**

10 The `omp_add_allocator_traits` routine adds an allocator trait to the allocator traits set.

11 **Format**

▼──────────────────────────────── C / C++ ─────────────────────────────────►

```
void omp_add_allocator_traits (omp_allocator_set_t *set,
                             size_t ntraits,
                             omp_allocator_t *traits);
```

▲──────────────────────────────── C / C++ ─────────────────────────────────▲

▼──────────────────────────────── Fortran ─────────────────────────────────►

```
subroutine omp_add_allocator_traits ( set, ntraits, traits )
integer(kind=omp_allocator_set_kind),intent(inout) :: set
integer,intent(in) :: ntraits
type(omp_allocator),intent(in) :: traits(*)
```

▲──────────────────────────────── Fortran ─────────────────────────────────▲

12 **Binding**

13 The binding thread set for an `omp_add_allocator` region is all threads on a device.
 14 The effect of executing this routine is not related to any specific region corresponding to
 15 any construct or API routine.

1 **Constraints on Arguments**

2 If the `ntraits` argument is greater than zero, then there must be at least as many traits
3 specified in the `traits` argument; otherwise, the behavior is unspecified.

4 **Effect**

5 The effect of the `omp_add_alloctrails` routine is that the `ntraits` specified in `traits` are
6 added to the set of allocator traits.

7 **Cross References**

- 8
 - [Allocator traits in Section 2.5.2 on page 52](#)

9 **3.5.11 omp_merge_alloctrails**10 **Summary**

11 The `omp_merge_alloctrails` routine merges two allocator traits sets.

12 **Format**

▼──────────────────────────────── C / C++ ─────────────────────────────────▶

```
void omp_merge_alloctrails (omp_alloctrail_set_t *dst,
                           const omp_alloctrail_set_t *src,
                           int dst_priority); (C)
void omp_merge_alloctrails (omp_alloctrail_set_t *dst,
                           const omp_alloctrail_set_t *src,
                           bool dst_priority=true); (C++)
```

▲──────────────────────────────── C / C++ ─────────────────────────────────▲

▼──────────────────────────────── Fortran ─────────────────────────────────▶

```
subroutine omp_merge_alloctrails ( dst, src, dst_priority )
integer(kind=omp_alloctrail_set_kind),intent(inout) :: dst
integer(kind=omp_alloctrail_set_kind),intent(in)  :: src
logical :: dst_priority
```

▲──────────────────────────────── Fortran ─────────────────────────────────▲

3.3.4 Routines for allocators

1 **Binding**

2 The binding thread set for an `omp_merge_alloctrails` region is all threads on a
3 device. The effect of executing this routine is not related to any specific region
4 corresponding to any construct or API routine.

5 **Effect**

6 The effect of the `omp_merge_alloctrails` routine is that the two allocator traits sets
7 `dst` and `src` are merged into `dst`. If the same memory trait appears in both sets, and the
8 `dst_priority` argument evaluates to `true` the merged value will be that of the `dst` set;
9 otherwise, it will be the value of the `src` set.

10 **3.5.12 `omp_init_allocator`**

11 **Summary**

12 The `omp_init_allocator` initializes an allocator and associates it with a memory
13 space.

14 **Format**

▼ C / C++ ▼

```
omp_allocator_t * omp_init_allocator ( omp_memspace_t *memspace,
                                     const omp_alloctrail_set_t *traits);
```

▲ C / C++ ▲

▼ Fortran ▼

```
integer(kind=omp_allocator_kind)
function omp_init_allocator ( memspace, traits )
integer(kind=omp_memspace_kind),intent(in) :: memspace
integer(kind=omp_alloctrail_set_kind),intent(in) :: traits
```

▲ Fortran ▲

1 **Binding**

2 The binding thread set for an `omp_init_allocator` region is all threads on a device.
 3 The effect of executing this routine is not related to any specific region corresponding to
 4 any construct or API routine.

5 **Constraints on Arguments**

6 The memspace argument must be a memory space returned by the
 7 `omp_init_memspace` routine. The traits argument must have been initialized with the
 8 `omp_init_alloctrait_set` routine.

9 **Effect**

10 The `omp_init_allocator` routine creates a new allocator that is associated with the
 11 memory space represented by the memspace handler. The allocations done through the
 12 created allocator will behave according to the allocator traits specified in the traits
 13 argument. Specifying the same allocator trait more than once results in unspecified
 14 behavior. The routine returns a handler for the created allocator. If the traits argument is
 15 an empty set this routine will always return a handler to an allocator. If the traits argument
 16 is not empty and an allocator that satisfies the requirements cannot be created then the
 17 special value `OMP_NULL_ALLOCATOR` is returned.

18 The traits in `omp_default_allocator_traits` must be defined as an empty set of
 19 allocator traits.

20 **Cross References**

- 21 • [Allocators in Section 2.5.2 on page 52](#)

22 **3.5.13 `omp_destroy_allocator`**

23 **Summary**

24 The `omp_destroy_allocator` releases all resources and memory allocations
 25 associated to an allocator.

26 **Format**

▼ `C / C++` ▼


```
void omp_destroy_allocator (omp_allocator_t *allocator);
```

▲────────────────── C / C++ ───────────────────▲

▼────────────────── Fortran ───────────────────▼

```
subroutine omp_destroy_allocator ( allocator )
integer(kind=omp_allocator_kind),intent(out) :: allocator
```

▲────────────────── Fortran ───────────────────▲

1 **Binding**

2 The binding thread set for an `omp_destroy_allocator` region is all threads on a
3 device. The effect of executing this routine is not related to any specific region
4 corresponding to any construct or API routine.

5 **Effect**

6 The `omp_destroy_allocator` routine releases resources that might be associated with
7 the allocator handler. Also, any memory allocated by the allocator but not deallocated yet
8 is deallocated by this routine.

9 **3.5.14 omp_set_default_allocator**

10 **Summary**

11 The `omp_set_default_allocator` sets the default allocator to be used by allocation
12 calls, directives and clauses that use default allocation.

13 **Format**

▼────────────────── C / C++ ───────────────────▼

```
void omp_set_default_allocator (omp_allocator_t *allocator);
```

▲────────────────── C / C++ ───────────────────▲

▼────────────────── Fortran ───────────────────▼

3.3.5 Routines for allocation/free

```
integer(kind=omp_allocator_kind)
function omp_get_default_allocator ()
```

Fortran

1 Binding

2 The binding task set for an `omp_get_default_allocator` region is the generating
3 task.

4 Effect

5 The effect of this routine is to return the value of the def-allocator-var ICV of the current
6 task.

7 Cross References

- 8 • def-allocator-var ICV, see Section 2.3 on page 39.
- 9 • `omp_alloc` routine, see Section 3.5.16 on page 347.

C / C++

10 3.5.16 `omp_alloc`

11 Summary

12 The `omp_alloc` requests a memory allocation to an allocator.

13 Format

```
void * omp_alloc (size_t size, omp_allocator_t *allocator); (C)
void * omp_alloc (size_t size,
                 omp_allocator_t *allocator=OMP_NULL_ALLOCATOR); (C++)
```

1 **Effect**

2 The `omp_alloc` routine requests a memory allocation of size bytes from the specified
 3 allocator without specifying an allocation alignment. If value of the allocator argument is
 4 `OMP_NULL_ALLOCATOR` the allocator used by the routine will be the one specified by the
 5 `def-allocator-var` ICV. Upon success it returns a pointer to the allocated memory.
 6 Otherwise, the behavior of the call depends on the `fallback` trait of the allocator.

7 **Cross References**

- 8 • [How Allocations Works](#), see Section 2.5.2 on page 52.

9 **3.5.17 omp_alloc_safe_align**10 **Summary**

11 The `omp_alloc_safe_align` requests a memory allocation to an allocator with an
 12 allocation alignment.

13 **Format**

```
void * omp_alloc_safe_align (size_t size, size_t alignment,
                             omp_allocator_t *allocator);           (C)
void * omp_alloc_safe_align (size_t size, size_t alignment,
                             omp_allocator_t *allocator=OMP_NULL_ALLOCATOR); (C++)
```

14 **Constraints on Arguments**

15 The allocator must be an allocator returned by the `omp_init_allocator` routine.
 16 Specifying an alignment argument that is not a power of two results in unspecified
 17 behavior.

18 **Effect**

19 The `omp_alloc_safe_align` routine requests a memory allocation of size bytes from
 20 the specified allocator where the allocation alignment of the request is alignment. If value
 21 of the allocator argument is `OMP_NULL_ALLOCATOR` the allocator used by the routine will
 22 be the one specified by the `def-allocator-var` ICV. Upon success it returns a pointer to the
 23 allocated memory. Otherwise, the behavior of the call depends on the `fallback` trait of
 24 the allocator.

1 **Cross References**

- 2 •
- [How Allocations Works](#)
- , see
- [Section 2.5.2](#)
- on page 52.

3 **3.5.18 omp_free**4 **Summary**5 The [omp_free](#) routine deallocates previously allocated memory.6 **Format**

```
void omp_free ( void * ptr, omp_allocator_t *allocator);       (C)
void omp_free ( void * ptr,
               omp_allocator_t *allocator = OMP_NULL_ALLOCATOR); (C++)
```

7 **Effect**

8 The [omp_free](#) routine deallocates the memory pointed by ptr. The ptr argument must
9 point to memory previously allocated with an OpenMP allocator. If the allocator is
10 specified it must be the allocator to which the allocation request was made. If the allocator
11 argument is [OMP_NULL_ALLOCATOR](#) the implementation will find the allocator used to
12 allocate the memory. Using [omp_free](#) on memory that was already deallocated results in
13 unspecified behavior.

▲────────────────── C / C++ ───────────────────▲

3.4 Changes to Chapter 5

1 **5.16 OMP_TOOL_LIBRARIES**

2 The **OMP_TOOL_LIBRARIES** environment variable sets the *tool-libraries-var* ICV to a list of tool
3 libraries that will be considered for use on a device where an OpenMP implementation is being
4 initialized. The value of this environment variable must be a comma-separated list of
5 dynamically-linked libraries, each specified by an absolute path.

6 If the *tool-var* ICV is not enabled, the value of *tool-libraries-var* will be ignored. Otherwise, if
7 **ompt_start_tool**, a global function symbol for a tool initializer, isn't visible in the address
8 space on a device where OpenMP is being initialized or if **ompt_start_tool** returns **NULL**, an
9 OpenMP implementation will consider libraries in the *tool-libraries-var* list in a left to right order.
10 The OpenMP implementation will search the list for a library that meets two criteria: it can be
11 dynamically loaded on the current device and it defines the symbol **ompt_start_tool**. If an
12 OpenMP implementation finds a suitable library, no further libraries in the list will be considered.

13 **Cross References**

- 14 • *tool-libraries-var* ICV, see Section 2.3 on page 39.
- 15 • Tool Interface, see Section 4 on page 364.
- 16 • **ompt_start_tool** routine, see Section 4.5.1 on page 396.

17 **5.17 OMP_ALLOCATOR**

18 The **OMP_ALLOCATOR** environment variable defines the memory and allocator traits to be
19 used to create the allocator to be set as the initial value of the *def-allocator-var* ICV.

20 The value of this environment variable is a comma-separated list of key=value elements
21 where each key is either a memory or allocator trait and value is one of the allowed values
22 for the specified trait.

23 **Cross References**

- 24 • memory and allocator traits, see Section 2.5 on page 50.
- 25 • *def-allocator-var* ICV, see Section 2.3 on page 39.
- 26 • **omp_set_default_allocator** routine, see Section 3.5.14 on page 345.
- 27 • **omp_get_default_allocator** routine, see Section 3.5.15 on page 346.

4 Examples

The examples presented in the section are intended to demonstrate how the proposed memory management APIs may be used in an OpenMP program. For each example, a C and Fortran version is presented. The example descriptions pertain to the C examples but apply to the corresponding Fortran examples unless otherwise noted. The first set of examples show how to use the APIs to perform dynamic memory allocation using default memory and allocator traits. The next set of examples demonstrate the APIs for explicitly specifying memory and allocator traits for dynamic memory allocation. The examples that follow show how variable declarations can be annotated with the declarative `allocate` directive. The final examples in this section show how allocation for private variables that arise from data-sharing clauses can be managed with the `allocate` clause.

4.1 Basic Allocation

First, we start with examples demonstrating how to use the memory management APIs to perform allocations with the default allocator. In the C example, `OMP_NULL_ALLOCATOR` is passed in to the `omp_alloc` call at line 10 indicating that the default allocator internally maintained by the implementation should be used. In the Fortran example, the same effect is achieved by annotating the `allocate` statement with an `allocate` directive without an `allocator` clause at line 8. Equivalently, the default allocator can be explicitly obtained and used in the code by using the `omp_get_default_allocator` routine.

The memory and allocator traits for the default allocator may be specified by using the `OMP_ALLOCATOR` environment variable or the `omp_set_default_allocator` routine; otherwise, its traits are implementation-defined. For example, suppose `OMP_ALLOCATOR` is set to “`optimized=bandwidth,fallback=abort_fb`” in the environment from which the program is executed and `omp_set_default_allocator` is not used. In this case, the allocation will occur from a bandwidth-optimized memory if it is available or else the program will abort.

C / C++

Example basic.1.c

```
S-1 #include <stdio.h>
S-2 #include <omp.h>
S-3
S-4 int basic_default1(int n)
S-5 {
S-6     const int success=1, failure=0;
S-7     int retval;
S-8     double *buffer;
S-9
S-10    buffer = omp_alloc(n * sizeof(*buffer), OMP_NULL_ALLOCATOR);
S-11
S-12    if (buffer == NULL) {
S-13        fprintf("Could not allocate space using default allocator\n");
S-14        retval = failure;
```

```

S-15     } else {
S-16         do_work(buffer, n);
S-17         omp_free(buffer, OMP_NULL_ALLOCATOR);
S-18         retval = success;
S-19     }
S-20
S-21     return retval;
S-22 }

```

▲────────────────────────────────── C / C++ ───────────────────────────────────▲

▼────────────────────────────────── Fortran ───────────────────────────────────▼

Example basic.1.f

```

S-1  function basic_default1(n) result(retval)
S-2      use omp_lib
S-3      integer :: n, retval
S-4      integer, parameter :: success=1, failure=0
S-5      double precision, allocatable :: buffer(:)
S-6      integer :: alloc_status
S-7
S-8      !$omp allocate
S-9      allocate(buffer(n), stat=alloc_status)
S-10
S-11     if (alloc_status /= 0) then
S-12         print *, "Could not allocate using default allocator"
S-13         retval = failure
S-14     else
S-15         call do_work(buffer, n)
S-16         deallocate(buffer)
S-17         retval = success
S-18     end if
S-19 end function basic_default1

```

▲────────────────────────────────── Fortran ───────────────────────────────────▲

The following examples shows how the proposed API can be used to perform memory allocation using default memory and allocator traits. The effect of using `omp_default_memtraits` is to request that the implementation assumes an implementation-defined set of default traits when selecting a memory for which a memory space object will be returned. The effect of using `omp_default_alloctraits` is to request that the implementation assumes the specified default values for each allocator trait when returning an allocator object, and it is therefore equivalent to setting up an allocator traits set object with zero added traits.

The call at line 12 is guaranteed to return a non-NULL value, and likewise the call at line 13 is guaranteed to return a non-NULL value. The resulting allocator may then be used for default allocations without any traits specified explicitly.

C / C++

Example basic.2.c

```

S-1 #include <stdio.h>
S-2 #include <omp.h>
S-3
S-4 int basic_default2(int n)
S-5 {
S-6     const int success=1, failure=0;
S-7     int retval;
S-8     omp_memspace_t *my_mspace;
S-9     omp_allocator_t *my_allocator;
S-10    double *buffer;
S-11
S-12    my_mspace = omp_init_memspace(&omp_default_memtraits);
S-13    my_allocator = omp_init_allocator(my_mspace, &omp_default_alloctrtraits);
S-14    buffer = omp_alloc(n * sizeof(*buffer), my_allocator);
S-15
S-16    if (buffer == NULL) {
S-17        fprintf("Could not allocate space using default traits\n");
S-18        retval = failure;
S-19    } else {
S-20        do_work(buffer, n);
S-21        omp_free(buffer, my_allocator);
S-22        retval = success;
S-23    }
S-24
S-25    omp_destroy_allocator(my_allocator);
S-26    omp_destroy_mspace(my_mspace);
S-27
S-28    return retval;
S-29 }

```

C / C++

Fortran

Example basic.2.f

```

S-1 function basic_default2(n) result(retval)
S-2     use omp_lib
S-3     integer :: n, retval
S-4     integer, parameter :: success=1, failure=0
S-5     integer (kind=omp_memspace_kind) :: my_mspace
S-6     integer (kind=omp_allocator_kind) :: my_allocator
S-7     double precision, allocatable :: buffer(:)
S-8     integer :: alloc_status
S-9
S-10    my_mspace = omp_init_memspace(omp_my_memtraits)

```



```
S-11     my_allocator = omp_init_allocator(my_mspace, omp_my_alloctraits)
S-12
S-13     !$omp allocate allocator(my_allocator)
S-14     allocate(buffer(n), stat=alloc_status)
S-15
S-16     if (alloc_status /= 0) then
S-17         print *, "Could not allocate using default traits"
S-18         retval = failure
S-19     else
S-20         call do_work(buffer, n)
S-21         deallocate(buffer)
S-22         retval = success
S-23     end if
S-24
S-25     call omp_destroy_allocator(my_allocator)
S-26     call omp_destroy_memspace(my_mspace)
S-27 end function basic_default2
```

Fortran

4.2 Allocation with Traits

In the following examples, the program attempts to allocate out of the memory providing the highest bandwidth while also supporting 2 megabyte pages. At lines 12 through 17, a memory space object is requested with the `bandwidth` trait set to `highest` and the `pagesize` trait set to 2 megabytes. Using the `bandwidth` trait rather than the `optimized` trait means that the memory providing the highest bandwidth while supporting 2MB pages should be used, regardless of whether it is actually designated as “bandwidth-optimized.” If the implementation is unable to return such a memory space object since a memory with a 2MB page size is unavailable, a memory space object with default traits is obtained. Next, the program requests an allocator object using the memory space object (pointed to by `mspace`) and default allocator traits.

The allocation is performed at line 19 using the obtained allocator. If the allocator is unable to allocate the requested number of bytes, then the implementation invokes the default fallback behavior – allocating, with default allocator traits, from a memory space with default memory traits. Even with this fallback behavior, it is possible that the allocation is ultimately unsuccessful. In this event the program returns from the function with a `failure` status.

C / C++

Example basic_traits.1.c

```
S-1 #include <stdio.h>
S-2 #include <omp.h>
S-3
S-4 int basic_traits1(int n)
S-5 {
S-6     const int success=1, failure=0;
S-7     const omp_memtrait_t mtrait_list[2] =
S-8         { {OMP_MTK_BANDWIDTH, OMP_MTV_HIGHEST},
S-9           {OMP_MTK_PAGESIZE, 2*1024*1024} };
S-10    int retval = success;
S-11
S-12    omp_memtrait_set_t mtraits;
S-13    omp_init_memtrait_set(&mtraits, 2, mtrait_list);
S-14    omp_memspace_t *my_mspace = omp_init_memspace(&mtraits);
S-15    if (my_mspace == OMP_NULL_MEMSPACE) {
S-16        my_mspace = omp_init_memspace(&omp_default_memtraits);
S-17    }
S-18
S-19    omp_allocator_t *my_allocator = omp_init_allocator(my_mspace,
S-20                                                       &omp_default_alloctrails);
S-21
S-22    double *buffer = omp_alloc(N * sizeof(*buffer), my_allocator);
S-23    if (buffer == NULL) {
S-24        fprintf(stderr, "Could not allocate using memory allocator\n");
S-25        retval = failure;
S-26    } else {
S-27        do_work(buffer, n);
S-28        omp_free(buffer, my_allocator);
S-29        retval = success;
S-30    }
S-31
S-32    omp_destroy_allocator(my_allocator);
S-33    omp_destroy_mspace(my_mspace);
S-34
S-35    return retval;
S-36 }
S-37
```

C / C++

 Fortran

Example basic_traits.1.f

```

S-1  function basic_traits1(n) result(retval)
S-2      use omp_lib
S-3      integer :: n, retval
S-4      integer, parameter :: success=1, failure=0
S-5      type(omp_memtrait), parameter :: mtrait_list(2) = &
S-6          (/ omp_memtrait(omp_mtk_bandwidth, omp_mtv_highest), &
S-7             omp_memtrait(omp_mtk_pagesize, 2*1024*1024) /)
S-8      integer (kind=omp_memtrait_set_kind) :: mtraits
S-9      integer (kind=omp_memspace_kind) :: my_mspace
S-10     integer (kind=omp_allocator_kind) :: my_allocator
S-11     double precision, allocatable :: buffer(:)
S-12     integer :: alloc_status
S-13
S-14     call omp_init_memtrait_set(mtraits, 2, mtrait_list)
S-15     my_mspace = omp_init_memspace(mtraits)
S-16     if (my_mspace == omp_null_memspace) then
S-17         my_mspace = omp_init_memspace(omp_default_mtraits)
S-18     end if
S-19
S-20     my_allocator = omp_init_allocator(my_mspace, omp_default_alloctrails)
S-21
S-22     !$omp allocate allocator(my_allocator)
S-23     allocate(buffer(n), stat=alloc_status)
S-24     if (alloc_status /= 0) then
S-25         print *, "Could not allocate using memory allocator"
S-26         retval = failure
S-27     else
S-28         call do_work(buffer, n)
S-29         deallocate(buffer)
S-30         retval = success
S-31     end if
S-32
S-33     call omp_destroy_allocator(my_allocator)
S-34     call omp_destroy_memspace(my_mspace)
S-35 end function basic_traits1

```

 Fortran

The next examples are similar to the previous ones, except here the program requires that the buffer is either allocated from a bandwidth-optimized (HBW) memory or returns from the function call with a `failure` status. At lines 19 through 22 the program explicitly requests an allocator having a fallback trait with the `null_fb` value. This means that if the allocator is unable to allocate the requested number of bytes at line 29 then a NULL value will be returned and the function will return with a `failure` status.

C / C++

Example basic_traits.2.c

```

S-1  #include <stdio.h>
S-2  #include <omp.h>
S-3
S-4  int basic_traits2(int n)
S-5  {
S-6      const int success=1, failure=0;
S-7      const omp_memtrait_t mtrait_list[1] =
S-8          { {OMP_MTK_OPTIMIZED, OMP_MTV_BANDWIDTH} };
S-9      omp_memtrait_set_t mtraits;
S-10     omp_init_memtrait_set(&mtraits, 1, mtrait_list);
S-11     omp_memspace_t *hbw_mspace = omp_init_memspace(&mtraits);
S-12     int retval;
S-13
S-14     if (hbw_mspace == OMP_NULL_MEMSPACE) {
S-15         fprintf(stderr, "Could not create memspace object for HBW memory\n");
S-16         retval = failure;
S-17     } else {
S-18         omp_alloctrait_set_t atrait;
S-19         const omp_alloctrait_t atrait_list[1] =
S-20             { {OMP_ATK_FALLBACK, OMP_ATV_NULL_FB} };
S-21         omp_init_alloctrait_set(&atrait, 1, atrait_list);
S-22         omp_allocator_t *hbw_allocator = omp_init_allocator(hbw_mspace, &atrait);
S-23
S-24         if (hbw_allocator == OMP_NULL_ALLOCATOR) {
S-25             fprintf(stderr, "Could not create allocator object for HBW memory\n");
S-26             retval = failure;
S-27         } else {
S-28
S-29             double *buffer = omp_alloc(N * sizeof(*buffer), hbw_allocator);
S-30             if (buffer == NULL) {
S-31                 fprintf(stderr, "Could not allocate using HBW memory allocator\n");
S-32                 retval = failure;
S-33             } else {
S-34                 do_work(buffer, n);
S-35                 omp_free(buffer, hbw_allocator);
S-36                 retval = success;
S-37             }
S-38             omp_destroy_allocator(hbw_allocator);
S-39         }
S-40         omp_destroy_memspace(hbw_mspace);
S-41     }
S-42
S-43     return retval;
S-44 }

```

S-45
S-46



Example basic_traits.2.f

```

S-1  function basic_traits2(n, result(retval)
S-2      use omp_lib
S-3      integer :: n, retval
S-4      integer, parameter :: success=1, failure=0
S-5      type(omp_memtrait), parameter :: mtrait_list(1) = &
S-6          (/ omp_memtrait(omp_mtk_optimized, omp_mtv_bandwidth) /)
S-7      integer (kind=omp_memtrait_set_kind) :: mtraits
S-8      integer (kind=omp_memspace_kind) :: hbw_mspace
S-9      type(omp_alloctrain), parameter :: atrait_list(1) = &
S-10         (/ omp_alloctrain(omp_atk_fallback, omp_atv_null_fb) /)
S-11     integer (kind=omp_alloctrain_set_kind) :: atraits
S-12     integer (kind=omp_allocator_kind) :: hbw_allocator
S-13     double precision, allocatable :: buffer(:)
S-14     integer :: alloc_status
S-15
S-16     call omp_init_memtrait_set(mtraits, 1, mtrait_list)
S-17     hbw_mspace = omp_init_memspace(mtraits)
S-18     if (hbw_mspace == omp_null_memspace) then
S-19         print *, "Could not create memspace object for HBW memory"
S-20         retval = failure
S-21     else
S-22         call omp_init_alloctrain_set(atraits, 1, atrait_list)
S-23         hbw_allocator = omp_init_allocator(hbw_mspace, atraits)
S-24
S-25         if (hbw_allocator == omp_null_allocator) then
S-26             print *, "Could not create allocator object for HBW memory"
S-27             retval = failure
S-28         else
S-29             !$omp allocate allocator(hbw_allocator)
S-30             allocate(buffer(n), stat=alloc_status)
S-31             if (alloc_status /= 0) then
S-32                 print *, "Could not allocate using memory allocator"
S-33                 retval = failure
S-34             else
S-35                 call do_work(buffer, n)
S-36                 deallocate(buffer)
S-37                 retval = success
S-38             end if
S-39             call omp_destroy_allocator(hbw_allocator)
S-40         end if

```

```

S-41         call omp_destroy_memspace(hbw_mspace)
S-42     end if
S-43 end function basic_traits2

```

▶────────────────── Fortran ─────────────────▶

4.3 Annotating Variable Declarations

In the following examples, a local array, `scratch`, is declared with length `n` and is used to perform local processing. Memory and allocator traits are explicitly specified on the `allocate` directive for `scratch`. The lifetime of the array is the duration of the call to `process_data`, as it would be if the `allocate` directive was not present. The implementation will therefore take care of performing the implicit deallocation of the array just prior to returning from the function.

◀────────────────── C / C++ ─────────────────▶

Example `allocate_directive.1.c`

```

S-1 #include <string.h>
S-2 #include <omp.h>
S-3
S-4 void process_data1(double *dat, size_t n)
S-5 {
S-6     double scratch[n];
S-7     #pragma omp allocate(scratch) memtraits(optimized=bandwidth) \
S-8         alloctraits(fallback=fb_abort)
S-9
S-10     memcpy(scratch, dat, n * sizeof(*dat));
S-11     do_local_work(scratch, n);
S-12     memcpy(dat, scratch, n * sizeof(*dat));
S-13 }

```

▶────────────────── C / C++ ─────────────────▶

 Fortran

Example allocate_directive.1.f

```

S-1  subroutine process_data1(dat, n)
S-2      use omp_lib
S-3      double precision :: dat(*)
S-4      integer :: n
S-5      double precision :: scratch(n)
S-6      !$omp allocate(scratch) memtraits(optimized=bandwidth) &
S-7      !$omp&                alloctrails(fallback=fb_abort)
S-8
S-9      scratch(1:n) = dat(1:n)
S-10     call do_local_work(scratch, n)
S-11     dat(1:n) = scratch(1:n)
S-12 end subroutine process_data2

```

 Fortran

In the next examples, again there is a local scratch array that is followed by an `allocate` directive. This time, an allocator object passed in as an argument is used to allocate `scratch`. The program requires that the local array be allocated in a bandwidth-optimized memory, and if it is unable to do so the program should abort.

 C / C++

Example allocate_directive.2.c

```

S-1  #include <string.h>
S-2  #include <omp.h>
S-3
S-4  void process_data2(double *dat, size_t n, omp_allocator_t *my_allocator)
S-5  {
S-6      double scratch[n];
S-7      #pragma omp allocate(scratch) allocator(my_allocator)
S-8
S-9      memcpy(scratch, dat, n * sizeof(*dat));
S-10     do_local_work(scratch, n);
S-11     memcpy(dat, scratch, n * sizeof(*dat));
S-12 }

```

 C / C++

 Fortran

Example allocate_directive.2.f

```

S-1  subroutine process_data2(dat, n, my_allocator)
S-2      use omp_lib
S-3      double precision :: dat(*)
S-4      integer :: n
S-5      integer (kind=omp_allocator_kind) :: my_allocator
S-6      double precision :: scratch(n)
S-7      !$omp allocate(scratch) allocator(my_allocator)
S-8
S-9      scratch(1:n) = dat(1:n)
S-10     call do_local_work(scratch, n)
S-11     dat(1:n) = scratch(1:n)
S-12 end subroutine process_data1

```

 Fortran

The next examples show how the `allocator_fb` fallback trait can be used. This time, a pointer to a structure containing user-defined allocators is passed in as an arguments. The `allocate` directive is used to allocate the local array in bandwidth-optimized memory, and if that is not possible it says the array should be allocated as per the allocator pointed to by `allocators->lat_opt`. The calling function, `process_data`, initializes the allocators with a call to `init_allocators` (line 32), and subsequently destroys the allocators with a call to `destroy_allocators` (line 34). It is also necessary to keep track of the memory space objects corresponding to each allocator since the lifetime of an allocator must not extend past the lifetime of its memory space.

 C / C++

Example allocate_directive.3.c

```

S-1  #include <string.h>
S-2  #include <omp.h>
S-3
S-4  struct allocators_t {
S-5      omp_allocator_t *bw_opt;
S-6      omp_allocator_t *lat_opt;
S-7      omp_allocator_t *cap_opt;
S-8      omp_memspace_t *bw_opt_mspace;
S-9      omp_memspace_t *lat_opt_mspace;
S-10     omp_memspace_t *cap_opt_mspace;
S-11 };
S-12
S-13 void process_data3(double *dat, size_t n, struct allocators_t *allocators)
S-14 {
S-15     double scratch[n];
S-16     #pragma omp allocate(scratch) memtraits(optimized=bandwidth) \
S-17         alloctrails(fallback=allocator_fb) \

```



```


S-18                                     alloctraits(fb_data=allocators->lat_opt)
S-19
S-20
S-21         memcpy(scratch, dat, n * sizeof(*dat));
S-22         do_local_work(scratch, n);
S-23         memcpy(dat, scratch, n * sizeof(*dat));
S-24     }
S-25
S-26 void init_allocators(struct allocators_t *allocators);
S-27 void destroy_allocators(struct allocators_t *allocators);
S-28
S-29 void process_data(double *dat, size_t n)
S-30 {
S-31     struct allocators_t allocators;
S-32     init_allocators(&allocators);
S-33     process_data3(dat, n, &allocators);
S-34     destroy_allocators(&allocators);
S-35 }
S-36
S-37 void init_allocators(struct allocators_t *allocators)
S-38 {
S-39     omp_memtrait_set_t mtraits;
S-40     omp_memtrait_t mtrait_list[1];
S-41
S-42     mtrait_list[0].key = OMP_MTK_OPTIMIZED;
S-43
S-44     /* create bandwidth-optimized allocator */
S-45     mtrait_list[0].value = OMP_MTV_BANDWIDTH;
S-46     omp_init_memtrait_set(&mtraits, 1, mtrait_list);
S-47     const omp_memspace_t *bw_opt_mspace = omp_init_memspace(&mtraits);
S-48     omp_destroy_memtrait_set(&mtraits);
S-49     allocators->bw_opt_mspace = bw_opt_mspace;
S-50     allocators->bw_opt = omp_init_allocator(bw_opt_mspace,
S-51                                           &omp_default_alloctraits);
S-52
S-53     /* create latency-optimized allocator */
S-54     mtrait_list[0].value = OMP_MTV_LATENCY;
S-55     omp_init_memtrait_set(&mtraits, 1, mtrait_list);
S-56     const omp_memspace_t *lat_opt_mspace = omp_init_memspace(&mtraits);
S-57     omp_destroy_memtrait_set(&mtraits);
S-58     allocators->lat_opt_mspace = lat_opt_mspace;
S-59     allocators->lat_opt = omp_init_allocator(lat_opt_mspace,
S-60                                           &omp_default_alloctraits);
S-61
S-62     /* create capacity-optimized allocator */
S-63     mtrait_list[0].value = OMP_MTV_CAPACITY;
S-64     omp_init_memtrait_set(&mtraits, 1, mtrait_list);

```

```

S-65     const omp_memspace_t *cap_opt_mspace = omp_init_memspace(&mtraits);
S-66     omp_destroy_memtrait_set(&mtraits);
S-67     allocators->cap_opt_mspace = cap_opt_mspace;
S-68     allocators->cap_opt = omp_init_allocator(cap_opt_mspace,
S-69                                           &omp_default_alloctrails);
S-70 }
S-71
S-72 void destroy_allocators(struct allocators_t *allocators)
S-73 {
S-74     omp_destroy_allocator(allocators->bw_opt);
S-75     omp_destroy_memspace(allocators->bw_opt_mspace);
S-76     omp_destroy_allocator(allocators->lat_opt);
S-77     omp_destroy_memspace(allocators->lat_opt_mspace);
S-78     omp_destroy_allocator(allocators->cap_opt);
S-79     omp_destroy_memspace(allocators->cap_opt_mspace);
S-80 }
S-81

```



Example allocate_directive.3.f

```

S-1 module mo_allocators
S-2     use omp_lib
S-3     type allocators_type
S-4         integer (omp_allocator_kind) :: bw_opt
S-5         integer (omp_allocator_kind) :: lat_opt
S-6         integer (omp_allocator_kind) :: cap_opt
S-7         integer (omp_memspace_kind)  :: bw_opt_mspace
S-8         integer (omp_memspace_kind)  :: lat_opt_mspace
S-9         integer (omp_memspace_kind)  :: cap_opt_mspace
S-10    end type
S-11 end module mo_allocators
S-12
S-13 subroutine process_data3(dat, n, allocators)
S-14     use mo_allocators
S-15     double precision :: dat(*)
S-16     integer :: n
S-17     type(allocators_type) :: allocators
S-18     double precision :: scratch(n)
S-19     !$omp allocate(scratch) memtraits(optimized=bandwidth) &
S-20     !$omp&                alloctrails(fallback=allocator_fb) &
S-21     !$omp&                alloctrails(fb_data=allocators%lat_opt)
S-22
S-23     scratch(1:n) = dat(1:n)
S-24     call do_local_work(scratch, n)
S-25     dat(1:n) = scratch(1:n)

```

```

S-26   end subroutine process_data3
S-27
S-28   subroutine init_allocators(allocators)
S-29       use omp_lib
S-30       use mo_allocators
S-31       type(allocators_type) :: allocators
S-32       integer (kind=omp_memtrait_set_kind) :: mtraits
S-33       type(omp_memtrait) :: mtrait_list(1)
S-34
S-35       mtrait_list(1)%key = omp_mtk_optimized
S-36
S-37       ! create bandwidth-optimized allocator
S-38       mtrait_list(1)%value = omp_mtv_bandwidth
S-39       call omp_init_memtrait_set(mtraits, 1, mtrait_list)
S-40       allocators%bw_opt_mspace = omp_init_memspace(mtraits)
S-41       call omp_destroy_memtrait_set(mtraits)
S-42       allocators%bw_opt = omp_init_allocator(allocators%bw_opt_mspace, &
S-43                                           omp_default_alloctrails)
S-44
S-45       ! create latency-optimized allocator
S-46       mtrait_list(1)%value = omp_mtv_latency
S-47       call omp_init_memtrait_set(mtraits, 1, mtrait_list)
S-48       allocators%lat_opt_mspace = omp_init_memspace(mtraits)
S-49       call omp_destroy_memtrait_set(mtraits)
S-50       allocators%lat_opt = omp_init_allocator(allocators%lat_opt_mspace, &
S-51                                           omp_default_alloctrails)
S-52
S-53       ! create capacity-optimized allocator
S-54       mtrait_list(1)%value = omp_mtv_capacity
S-55       call omp_init_memtrait_set(mtraits, 1, mtrait_list)
S-56       allocators%cap_opt_mspace = omp_init_memspace(mtraits)
S-57       call omp_destroy_memtrait_set(mtraits)
S-58       allocators%cap_opt = omp_init_allocator(allocators%cap_opt_mspace, &
S-59                                           omp_default_alloctrails)
S-60   end subroutine init_allocators
S-61
S-62   subroutine destroy_allocators(allocators)
S-63       use mo_allocators
S-64       type(allocators_type) :: allocators
S-65
S-66       call omp_destroy_allocator(allocators%bw_opt)
S-67       call omp_destroy_memspace(allocators%bw_opt_mspace)
S-68       call omp_destroy_allocator(allocators%lat_opt)
S-69       call omp_destroy_memspace(allocators%lat_opt_mspace)
S-70       call omp_destroy_allocator(allocators%cap_opt)
S-71       call omp_destroy_memspace(allocators%cap_opt_mspace)
S-72   end subroutine destroy_allocators

```

```
S-73
S-74 subroutine process_data(dat, n)
S-75     use mo_allocators
S-76     double precision :: dat(*)
S-77     integer :: n
S-78     type(allocators_type) :: allocators
S-79
S-80     call init_allocators(allocators)
S-81     call process_data3(dat, n, allocators)
S-82     call destroy_allocators(allocators)
S-83 end subroutine process_data
```

Fortran

4.4 Memory Management for Privatized Variables

The following examples illustrate the use of the `allocate` clause. A parallel loop is used to perform an array reduction across rows of a 2-dimensional array, `b`, which has been allocated in bandwidth-optimized memory. Each private copy of the 1-dimensional array, `a`, resulting from the `reduction` clause is allocated according to the `allocate` clause. In this case, the program requests that each thread's private array is also allocated in bandwidth-optimized memory.

C / C++


Example `allocate_clause.1.c`

```
S-1 #include <stdio.h>
S-2
S-3 #define N 100
S-4 void init(int n, float (*b)[N]);
S-5
S-6 int main()
S-7 {
S-8     int i, j;
S-9     float a[N], b[N][N];
S-10     #pragma allocate(a, b) memtraits(optimized=bandwidth)
S-11
S-12     init(N, b);
S-13
S-14     for (i = 0; i < N; i++) a[i] = 0.0e0;
S-15
S-16     #pragma omp parallel for reduction(+:a) private(j) \
S-17         allocate(memtraits(optimized=bandwidth):a)
S-18     for (i = 0; i < N; i++) {
S-19         for (j = 0; j < N; j++) {
S-20             a[j] += b[i][j];
S-21         }
    }
```

```

S-22     }
S-23
S-24     printf(" a[0] a[N-1]: %f %f \n", a[0], a[N-1]);
S-25
S-26     return 0;
S-27 }
S-28

```



Example allocate_clause.1.f

```

S-1  program array_red
S-2      integer, parameter  :: n=100
S-3      integer            :: j
S-4      real               :: a(n), b(n,n)
S-5      !$omp allocate(a, b) memtraits(optimized=bandwidth)
S-6
S-7      call init(n,b)
S-8
S-9      a(:) = 0.0e0
S-10
S-11     !$omp parallel do reduction(+:a) allocate(memtraits(optimized=bandwidth):a)
S-12     do j = 1, n
S-13         a(:) = a(:) + b(:,j)
S-14     end do
S-15
S-16     print *, " a(1) a(n): ", a(1), a(n)
S-17 end program

```



5 Next steps

This document outlines multiple additions to the OpenMP specification to augment it with an initial modern memory management interface that is capable of supporting the new and future memory technologies but we believe that more features are needed to fully cover all programmer needs. The following are the areas, in no particular order, in which we expect to continue to work targeting the future OpenMP 5.0 specification:

- **Host-device interaction.** The presented mechanisms can be used from within a target region to manage the device memory but do not allow to manage it from the host device. We envision two extensions in this direction:
 1. Allow the **allocate** clause to appear in target directives to affect the device allocations that arise from the **map** clauses.
 2. Extend the API to allow creation of device allocator and allocating memory using these allocators in a similarly to the existing **omp_target_alloc** routine.
- **Predefined trait sets.** We plan to provide a set of standard defined trait sets that encode requirements (e.g., high-bandwidth memory or scratchpad memories) and simplify for common cases of the API usage.
- **NUMA support.** We are exploring mechanisms that allow to distribute memory allocations across the different NUMA domains that could exist in a memory space.
- **Resource querying.** To enable maximum flexibility in looking for the appropriate memory spaces, we plan to develop an API that will allow to query which memory spaces exist in a system (and its attached devices) and which are the traits of each memory space.
- **C++ support.** We acknowledge that the current interface might not blend well with the usages of many C++ programmers and we intend to study how to improve this by providing either additional APIs that work with C++ types such as *std::vector* or redefined C++ operators and allocators.
- **Special code generation support.** Some existing and future memories require compilers to generate different code than for regular memories. Additional directives will be provided to guide the compiler in this process and to allow multiple versions of the same code to exist to work with different memories as necessary.
- **Static allocators.** In some cases in the current proposal we require users to provide an explicit list of traits instead of an allocator. This can get cumbersome and it goes against our principle of moving the decision away from the allocation place as the traits need to be repeated in each allocate directive or clause. To help overcome this problem we envision the ability to fully define allocators at compile time which can be used in places where a dynamic decision is not possible.