

# | GCC/OpenMP Update

Tobias Burnus  
Catherine Moore

# | Agenda

- **GCC Overview**
  - GCC Community
  - GCC Release Cycle
  - GCC 12 OpenMP Support
  - Testing OpenMP @Siemens
  - GCC Resources
- **OpenMP 5.0 + 5.1 Support**
- **Specification Corner-Cases**
- **Command-Line Options & Tricks**
- **Conclusion**

# OpenMP Community in GCC

- **Welcoming to new contributors**
  - Developer Certificate of Origin (DCO) or FSF copyright needed
- **Siemens (funded by ORNL and the DOE)**
  - Six active developers, led by Tobias Burnus, working on OpenMP functionality and performance
  - Major contributor; most of Fortran development
- **OpenMP Patch Review**
  - Maintainer - Jakub Jelinek
  - Reviews and contributes patches
- **Others**
  - Participation on ad-hoc basis

# GCC Release Cycle

Date	Release	Development Branch
April 2021	GCC 11	OG11 Branch
Spring 2022	GCC 12	OG12 Branch
Spring 2023	GCC 13	OG13 Branch

## OG Development Branches

- GCC GIT branch devel/omp/gcc-11 etc.
  - Maintained by Siemens developers
- Offers early access to OpenMP offloading features not in the official release
- Allows development to continue during GCC pre-release quiet periods
- Recommended for use for latest performance and functionality
- GCC Open Development: Usually May through November

# GCC 12 OpenMP Support

OpenMP Revision	Support Level	NVIDIA Offloading	AMD Offloading
OpenMP 4.5	Fully Supported	Yes	Yes
OpenMP 5.0	Partial Support	Yes	Yes
OpenMP 5.1	Minimal Support	Yes	Yes

AMD Offloading Support for MI100 is complete; Newer AMD offerings planned for 2022

# Testing OpenMP @Siemens

Nightly and weekly tests targeting AMD M100 and NVIDIA Volta

## Test suites

- solve\_vv
- omptest
- OvO
- Babelstream
  
- SPEC ACCEL
- SPECChpc 2021
  
- GCC's DejaGNU test suite
  - C/C++
  - gfortran
  - libgomp

# GCC Resources

GCC Offloading: <https://gcc.gnu.org/wiki/Offloading>

- Building GCC for offloading
- Options for building applications for offloading

OpenMP implementation status (development branch):

- <https://gcc.gnu.org/onlinedocs/libgomp/OpenMP-Implementation-Status.html>

Libgomp manual:

- GCC 11.2 (stable): <https://gcc.gnu.org/onlinedocs/gcc-11.2.0/libgomp/>
- Development branch: <https://gcc.gnu.org/onlinedocs/libgomp/>

General help with GCC

- mailing list: <https://gcc.gnu.org/pipermail/gcc-help/>

# | Agenda

- **GCC Overview**
- **OpenMP 5.0 + 5.1 Support**
  - **OpenMP Support in GCC**
  - **OpenMP 5.0 + 5.1 Features Supported in GCC 12**
  - **OpenMP 5.0 + 5.1 Features Unsupported, Planed for GCC 13**
- **Specification Corner-Cases**
- **Command-Line Options & Tricks**



# OpenMP Support in GCC

- GCC: Compiler for C, C++, Fortran, Ada, D, go, ...
  - C17 (steps to C2x), C++20 (steps to C++23)
  - Fortran 2008 + coarray + interop TS, initial F2018
- OpenMP and OpenACC 2.6 with C/C++/Fortran
- Supported archs: aarch64, alpha, arc, arm, avr, bfin, ...
- GCC offloading-support packages of Linux distributions
  - Debian/Ubuntu: gcc-11-offload-{nvptx,amdgc}
  - (open)SUSE: cross-{nvptx,amdgc}-gcc11
  - Red Hat/Fedora: {gcc,libgomp}-offload-nvptx (currently no amdgc)

## New features

- GCC 9 (2019): OpenMP 4.5 (C/C++, Fortran mostly), some 5.0, OpenACC 2.5
- GCC 10 (2020): More of OpenMP 5.0, OpenACC 2.6
- GCC 11 (2021): More of OpenMP 5.0
- GCC 12 (2022): Some OpenMP 5.1, more 5.0 (esp. Fortran)



### Supported Releases

#### **GCC 11.2** (changes)

Status: **2021-07-28** (regression fixes & docs only).  
Serious regressions. All regressions.

#### **GCC 10.3** (changes)

Status: **2021-04-08** (regression fixes & docs only).  
Serious regressions. All regressions.

#### **GCC 9.4** (changes)

Status: **2021-06-01** (regression fixes & docs only).  
Serious regressions. All regressions.

#### **Development:** GCC 12.0 (release criteria, changes)

Status: **2022-01-17** (regression fixes & docs only).  
Serious regressions. All regressions.

## OpenMP 5.0 Features Supported in GCC 12

Iterators	<b>omp_fulfill_event</b> runtime routine
<i>target-offload-var</i> ICV and OMP_TARGET_OFFLOAD env variable	<b>reduction</b> and <b>in_reduction</b> clauses on <b>taskloop</b> and <b>taskloop simd</b> constructs
Nested-parallel changes to <i>max-active-levels-var</i> ICV	<b>taskloop</b> construct cancelable by <b>cancel</b> construct
<b>teams</b> construct outside an enclosing target region	<b>mutexinuset</b> <i>dependence-type</i> for <b>depend</b> clause
<b>!=</b> as relational-op in canonical loop form for C/C++	Predefined memory spaces, memory allocators, allocator traits
<b>nonmonotonic</b> as default loop schedule modifier for worksharing-loop constructs	Memory management routines
Clauses <b>if</b> , <b>nontemporal</b> and <b>order(concurrent)</b> in <b>simd</b> construct	<b>use_device_addr</b> clause on target data
<b>atomic</b> constructs in <b>simd</b>	Implicit <b>declare target</b> directive
<b>loop</b> construct	C/C++'s lvalue expressions in depend clauses
<b>order(concurrent)</b> clause	Nested <b>declare target</b> directive
<b>scan</b> directive and <b>in_scan</b> modifier for the <b>reduction</b> clause	Combined <b>master</b> constructs
<b>in_reduction</b> clause on <b>task</b> constructs	<b>depend</b> clause on taskwait
<b>task_reduction</b> clause with <b>taskgroup</b>	Weak memory ordering clauses on <b>atomic</b> and <b>flush</b> construct
<b>task</b> modifier to <b>reduction</b> clause	<b>depobj</b> construct and depend objects
<b>affinity</b> clause to <b>task</b> construct	Lock hints were renamed to synchronization hints
<b>detach</b> clause to <b>task</b> construct	<b>conditional</b> modifier to <b>lastprivate</b> clause
<b>close</b> <i>map-type-modifier</i>	<b>defaultmap</b> extensions
<b>omp_get_supported_active_levels</b> routine	Runtime routines and environment variables to display runtime thread affinity information
<b>omp_pause_resource</b> and <b>omp_pause_resource_all</b> runtime routines	<b>omp_get_device_num</b> runtime routine
Supporting C++'s range-based for loop	

## OpenMP 5.1 Features Supported in GCC 12

OpenMP directive as C++ attribute specifiers

**nothing** directive

**error** directive

**masked** construct

**scope** directive

**strict** modifier in the **grainsize** and **num\_tasks** clauses of the **taskloop** construct

**thread\_limit** clause to **target** construct

Extensions to the **atomic** directive

**seq\_cst** clause on a **flush** construct

**private** and **firstprivate** argument to **default** clause in C and C++

**omp\_set\_num\_teams**, **omp\_set\_teams\_thread\_limit** runtime routines

**omp\_get\_max\_teams**, **omp\_get\_teams\_thread\_limit** runtime routines

**omp\_calloc**, **omp\_realloc** runtime routines

**omp\_aligned\_alloc** and **omp\_aligned\_calloc** runtime routines

**omp\_alloctrail\_key\_t** enum: **omp\_atv\_serialized** added, **omp\_atv\_default** changed

OMP\_PLACES syntax extensions

OMP\_NUM\_TEAMS and OMP\_TEAMS\_THREAD\_LIMIT environment variables

Support of strictly structured blocks in Fortran

Support of structured block sequences in C/C++

**unconstrained** and **reproducible** modifiers on order clause

**omp\_display\_env** runtime routine

## OpenMP 5.0 Features Unsupported in GCC 12

Feature	Development Branch Support	Planned for GCC 13
Array shaping	No	Yes
Array sections with non-unit strides in C and C++	No	Yes
<b>metadirective</b> directive	Yes	Yes
Collapse of associated loops that are imperfectly nested loops	No	Yes
<b>allocate</b> directive	Yes	Yes
Discontiguous array section with <b>target update</b> construct	No	Yes
C/C++'s lvalue expressions in <b>to</b> , <b>from</b> , and <b>map</b> clauses	No	Yes
<b>declare mapper</b> directive	No	Yes
OMPT interface	No	No
OMPD interface	No	No

## OpenMP 5.1 Features Planned for GCC 13 – and beyond

<b>omp_all_memory</b> reserved locator	
<i>target_device trait</i> in OpenMP Context	<b>omp_target_memcpy_async</b> runtime routine
<b>target_device</b> selector set in context selectors	<b>omp_target_memcpy_rect_async</b> runtime routine
C/C++'s <b>declare variant</b> directive: elision support of preprocessed code	<b>omp_get_mapped_ptr</b> runtime routine
<b>declare variant</b> : new clauses <b>adjust_args</b> and <b>append_args</b>	<i>ompt_scope_endpoint_t enum: ompt_scope_beginend</i>
<b>dispatch</b> construct	<i>ompt_sync_region_t enum additions</i>
device-specific ICV settings the environment variables	<i>ompt_state_t enum: ompt_state_wait_barrier_implementation</i>
<b>assume</b> directive	<i>ompt_state_t enum: ompt_state_wait_barrier_teams</i>
Loop transformation constructs	<i>ompt_callback_target_data_op_emi_t, ompt_callback_target_emi_t</i>
<b>has_device_addr</b> clause to target construct	<i>ompt_callback_target_map_emi_t and ompt_callback_target_submit_emi_t</i>
iterators in <b>target update</b> motion clauses and <b>map</b> clauses	<i>ompt_callback_error_t type</i>
indirect calls to the device version of a procedure or function in target regions	<b>nowait</b> clause in taskwait directive
<b>inoutset</b> argument to the <b>depend</b> clause	<b>interop</b> directive
<b>present</b> argument to <b>defaultmap</b> clause	<b>omp_interop_t</b> object support in runtime routines
<b>omp_target_is_accessible</b> runtime routine	

## OpenMP 5.0 Features Partially Supported in GCC 12

**declare variant** directive – *simd traits not handled correctly*

**requires** directive – *only fulfillable requirement are **atomic\_default\_mem\_order** and **dynamic\_allocators***

Non-rectangular loop nests – *Fortran support missing, planned for GCC 13*

**in\_reduction** clause on target constructs – *nowait is only a stub*

**ancestor** modifier on device clause – *reverse offload unsupported*

Map-order clarifications

Mapping C/C++ pointer variables and to assign the address of device memory mapped by an array section

Mapping of Fortran pointer and allocatable variables, including pointer and allocatable components of variables  
– *Mapping of vars with allocatable components unsupported, planned for GCC 13*

## OpenMP 5.1 Features Partially Supported in GCC 12

**align** clause/modifier in **allocate** directive/clause and **allocator** directive — *Fortran support missing, planned for GCC 13*

Planned for GCC 13:

- Unified shared memory support with NVIDIA GPUs

# | Agenda

- **GCC Overview**
- **OpenMP 5.0 + 5.1 Support**
- **Specification Corner Cases**
  - **OpenMP Specification**
  - **Example for a Minor Spec Issue**
  - **OMP\_TARGET\_OFFLOAD=mandatory**
  - **Spec Work from Our Side**
- **Command-Line Options & Tricks**
- **Conclusion**

# OpenMP Specification

## Roughly annual releases

- Either a new OpenMP release or a Technical Report (TR) as snapshot/preview of the next release
- Dot releases have minor changes
- Main change in 5.2: reorganization and syntax representation

## Stakeholders

- OpenMP Architecture Review Board (ARB) has 33 members representing GCC: SIEMENS, SUSE, Red Hat (via IBM)
- Weekly language spec meeting (plus subcommittees meetings)

## Issues keep popping up: unclear/underspecified, oversights, missing updates after changes/extensions elsewhere

- Despite: text-change discussions, two-step voting, pre-merge proof reading, and whole-document proof reading



# OpenMP Specification – Example of a Minor Spec Issue

## Fortran – Optional *End-Directive* and *Strictly Structured Blocks*

- Question: What ends the ‘end’
  - !\$omp parallel ! *Loosely structured block* – no ‘block’ follows
  - !\$omp parallel ! *Strictly structured block* – ‘block’ next
- block ! → ‘!\$omp end parallel’ is optional
- x = x + 1
- end block
- !\$omp end parallel ! ← ends outer or inner ‘parallel’?
  
- OpenMP 5.1: ?
- OpenMP 5.2 added: ‘An *end-directive* that immediately follows a directive ... is always paired with that directive.’
  - Applies to inner
  - An additional ‘!\$omp end parallel’ is missing

## OMP\_OFFLOAD\_TARGET=mandatory (1/4) When to Actually to Abort with an Error

### OpenMP has (5.0 + 5.1)

“The **mandatory** value specifies that program execution is terminated if a device construct or device memory routine is encountered and the device is not available or is not supported by the implementation.”

### Real-world situation

- System with installed GPU, CUDA installed but no device available due to kernel issue (cuInit: no CUDA-capable device is detected)  
→ `omp_get_num_devices() == 0` (→ only host)

### Question: Should fail with “mandatory” or not?

- GCC → host fallback/no fail – as no device exists, default device is the host
- LLVM → fails with the CUDA error
- User expectation: Hardware exists but does not work → should fail  
And if no hardware exists → no fail or still a fail?

# OMP\_OFFLOAD\_TARGET=mandatory (1/4)

## When to Actually to Abort with an Error

### OpenMP has (5.0 + 5.1)

“The **mandatory** value specifies that program execution is terminated if a device construct or device memory routine is encountered and the device is not available or is not supported by the implementation.”

### Real-world situation

- System with installed GPU, CUDA installed but no device available due to kernel issue (cuInit: no CUDA-capable device is detected)  
→ `omp_get_num_devices() == 0` (→ only host)

### Question: Should fail with “mandatory” or not?

- GCC → host fallback/no fail – as no device exists, default device is the host
- LLVM → fails with the CUDA error
- User expectation: Hardware exists but does not work → should fail  
And if no hardware exists → no fail or still a fail?

### Workaround

Environment variable: ?

*OMP\_DEFAULT\_DEVICE=kind(gpu)  
would be useful, but does not exist  
in OpenMP ≤ 5.2*

Code:

```
if (omp_get_default_device()  
    == omp_get_initial_device())  
    → error
```

## OMP\_OFFLOAD\_TARGET=mandatory (2/4)

### When to Actually to Abort with an Error

#### Question: What if no non-host device is available and ...

```
void foo () {}  
#pragma omp declare target to (foo)  
  
int main () {  
    #pragma omp target if(false)  
    foo ();      // Is this ok?  
    omp_set_default_device (omp_get_initial_device ());  
    #pragma omp target  
    foo ();      // What about this?  
    #pragma omp target device(omp_get_initial_device ())  
    foo ();      // Or this?  
    #pragma omp target device(omp_get_num_devices () + 42)  
    foo ();      // This one is clearly an error  
    if (omp_get_num_devices () == 3)  
    {  
        #pragma omp target device (1)  
        foo ();  // This would be an error if we can't offload to device 1  
    }  
}
```

## OMP\_OFFLOAD\_TARGET=mandatory (3/4) When to Actually to Abort with an Error

### Solution in OpenMP 5.2

- **Definitions** — “the constant **omp\_initial\_device** can be used as an alias for the host device and the constant **omp\_invalid\_device** can be used to specify an invalid device number. A *conforming device number* is either a non-negative integer that is less than or equal to **omp\_get\_num\_devices()** or equal to **omp\_initial\_device** or **omp\_invalid\_device**.”
- **default-device-var initialization** — “If *target-offload-var* is **mandatory** and the number of non-host devices is zero then the *default-device-var* is initialized to **omp\_invalid\_device**. Otherwise, the initial value is an implementation-defined non-negative integer that is less than or, if *target-offload-var* is not mandatory, equal to **omp\_get\_initial\_device()**.”
- **OMP\_OFFLOAD\_TARGET=mandatory** — “The **mandatory** value specifies that the effect of any device construct or device memory routine that uses a device that is unavailable or not supported by the implementation, or uses a non-conforming device number, is as if the **omp\_invalid\_device** device number was used.”

## OMP\_OFFLOAD\_TARGET=mandatory (4/4)

### When to Actually to Abort with an Error

#### Solution in OpenMP 5.2 (con'd)

- **device\_num** clause – “If the device-description evaluates to **omp\_invalid\_device**, runtime error termination is performed.”
- **Device Memory Routines:** “If the *device\_num*, *src\_device\_num*, or *dst\_device\_num* argument of a device memory routine has the value **omp\_invalid\_device**, runtime error termination is performed.”
- **Definition:** “When *runtime error termination* is performed, the effect is as if an **error** directive for which *sev-level* is **fatal** and *action-time* is **execution** is encountered.”

## Spec Work from Our Side

### Mainly trying to fix issues and improve wording

- Usually found when implementing a feature and looking at the spec
- Sometimes found by chance or forwarding issues reported to us
- Taking care of issues found during the (sub)committee discussion

### Recent examples (on going)

- Issues in Fortran part related to 5.1-added conditions support in 'atomic'
- has\_device\_addr – trying to clarify semantics
- Extend OMP\_DEFAULT\_DEVICE (cf. mandatory discussion)

### If you find a potential issue, bug, missing feature

- Check newer version of the spec – could be a bug which was fixed
- Contact some ARB member to take care of the bug (or participate if your org is already a member – or asked your org to become a member)

# | Agenda

- **GCC Overview**
- **OpenMP 5.0 + 5.1 Support**
- **Specification Corner Cases**
- **Command-line Options & Tricks**
  - **GCC Command Line Options for Offloading**
  - **Tricks & Tips – nvptx**
  - **Tricks & Tips – GCN**
- **Conclusion**



# GCC Command Line Options for Offloading

## Solution in OpenMP 5.2 (con'd)

- Generates offload code by default for OpenMP target regions with **-fopenmp** (as configured; for instance for both nvptx and AMD GCN)
- **-foffload=<disable|default|*target-list*>** – restrict to those device types
- **-foffload-options=<option>** or **-foffload-options=<target>=<options>**  
specify option to the offload compiler  
Typical examples:
  - foffload-options=-lgfortran -foffload-options=-lm
  - foffload-options="-lgfortran -lm" -foffload-options=nvptx-none=-latomic
  - foffload-options=amdgcn-amdhsa=-march=gfx906 -foffload-options=-lm
- Verbose optimization pass diagnostic: **-fopt-info[-options[=filename]]**  
Example: -foffload-options=-fopt-info-loop-missed -fopt-info-omp-missed

## Tricks & Tips

### nvptx

#### JIT

GCC generates generic code, which is just-in-time compiled by CUDA at startup – and cached in the user's directory.

→ <https://developer.nvidia.com/blog/cuda-pro-tip-understand-fat-binaries-jit-caching/>

→ `CUDA_CACHE_{DISABLE,MAXSIZE,PATH}`

#### GCC compile flags for nvptx

Usually not needed due to JIT

<https://gcc.gnu.org/onlinedocs/gcc/Nvidia-PTX-Options.html>

Possible exceptions:

- `-mptx=N.N` (PTX ISA version), `-misa=sm_XX`
- “illegal memory access was encountered” – generic error; could be stack issue, if so: `-foffload=-msoft-stack-reserve-local=...` might help.

Default 128 byte (note: multiplied by `sm_count × thread_max` ~ 20000)

## Tricks & Tips

### GCN

#### Hardware Specific Compilation

Native code for the specified GCN hardware is generated.

Use, e.g., `-foffload-options=-march=fiji` (GCN3, gfx803 – the default) or for GCN5 GPUs: `gfx900` (VEGA 10), `gfx906` (VEGA 20), or `gfx908`.

#### ROCGDB

Offloading debugging is supported with ROCGDB.

→ Slides: <https://linuxplumbersconf.org/event/11/contributions/997/>

→ Video via ↑ or <https://webinars.sw.siemens.com/debugging-offloaded-kernels-on-amd>

# Conclusion

## OpenMP in GCC

- OpenMP 5.0 mostly supported + initial 5.1 support in GCC 12

## Planned and/or useful

- GCC 13: Most of OpenMP 5.0, more of 5.1
- Improve device support: unified shared memory, performance, ...
- Diagnostic, documentation improvements

## Specification

- Large – and the devil is in the details

## Community effort

- Both the spec and the compiler depends on feedback, support, and work of users, developers (paid and hobby), and vendors

## Acknowledgement

This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725

## Disclaimer

© Siemens 2022

Subject to changes and errors. The information given in this document only contains general descriptions and/or performance features which may not always specifically reflect those described, or which may undergo modification in the course of further development of the products. The requested performance features are binding only when they are expressly agreed upon in the concluded contract.

All product designations may be trademarks or other rights of Siemens AG, its affiliated companies or other companies whose use by third parties for their own purposes could violate the rights of the respective owner.