



# AOMP: OpenMP Target Offloading for AMD GPUs

Jan-Patrick Lehr  
Member of Technical Staff  
Software Development Engineer

# Disclaimer

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. Any computer system has risks of security vulnerabilities that cannot be completely prevented or mitigated. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

THIS INFORMATION IS PROVIDED 'AS IS.' AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS, OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION. AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY RELIANCE, DIRECT, INDIRECT, SPECIAL, OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

© 2023 Advanced Micro Devices, Inc. All rights reserved.


AMD, the AMD Arrow logo, EPYC, Instinct and combinations thereof are trademarks of Advanced Micro Devices, Inc. PCIe is a registered trademark of PCI-SIG Corporation. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

---

# Agenda

- 
1. A high-level overview of the ROCm™ software stack
  2. AOMP and its (software) dependencies
  3. AOMP compilation/linking process
  4. AOMP architecture
  5. Example OpenMP target offload

# ROCm™ Software Stack



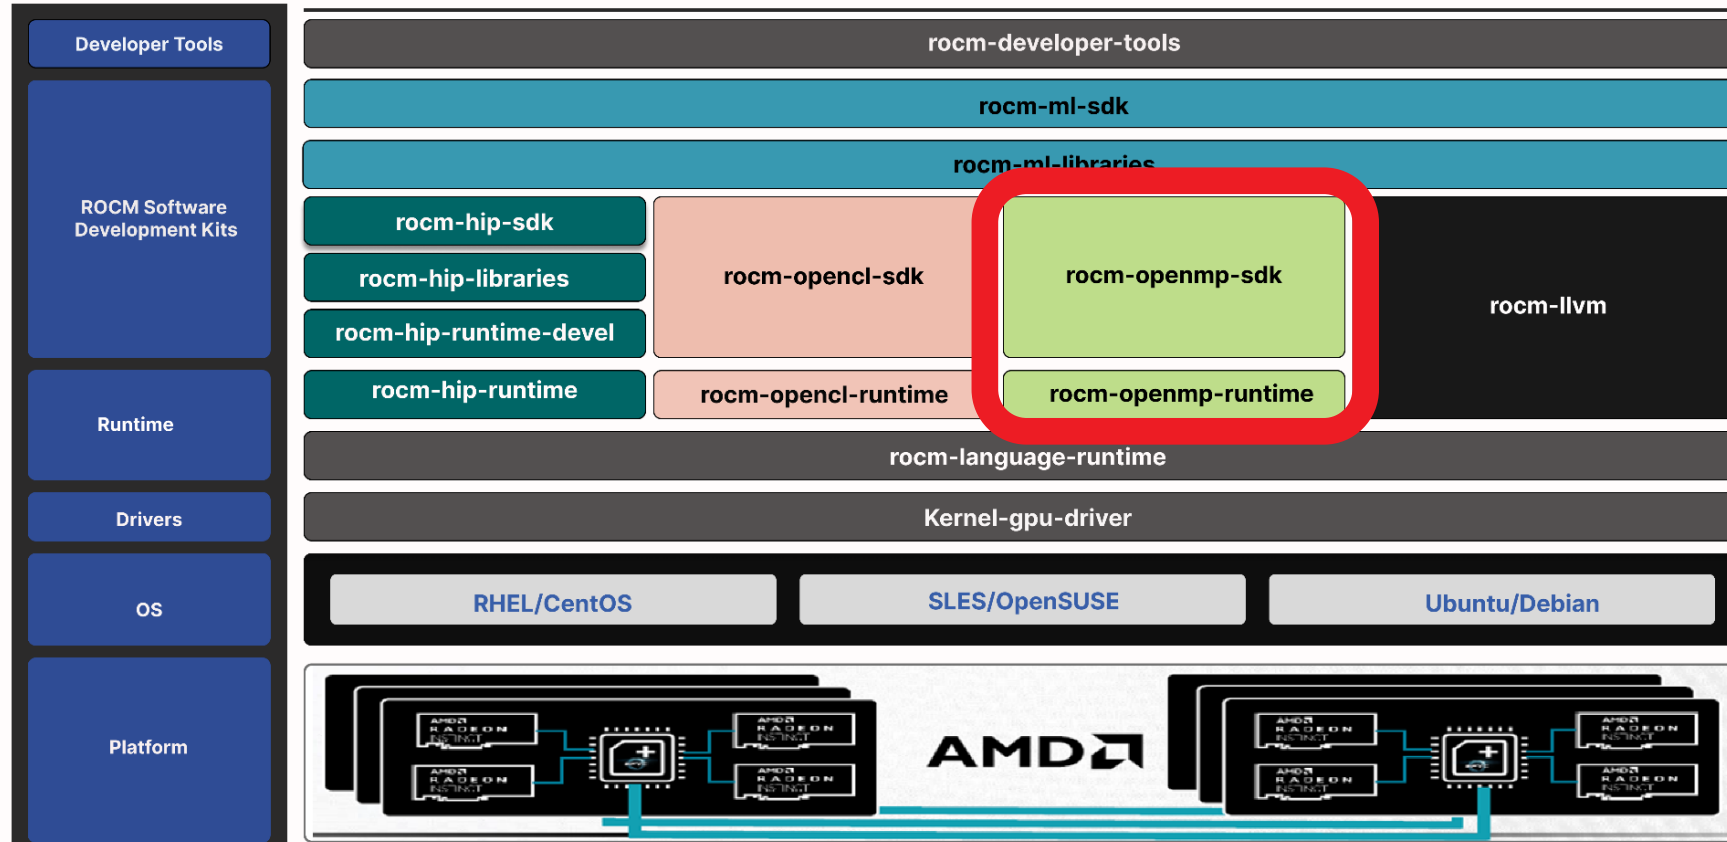


AMD  
**ROCm**

Open Software Platform for GPU Compute

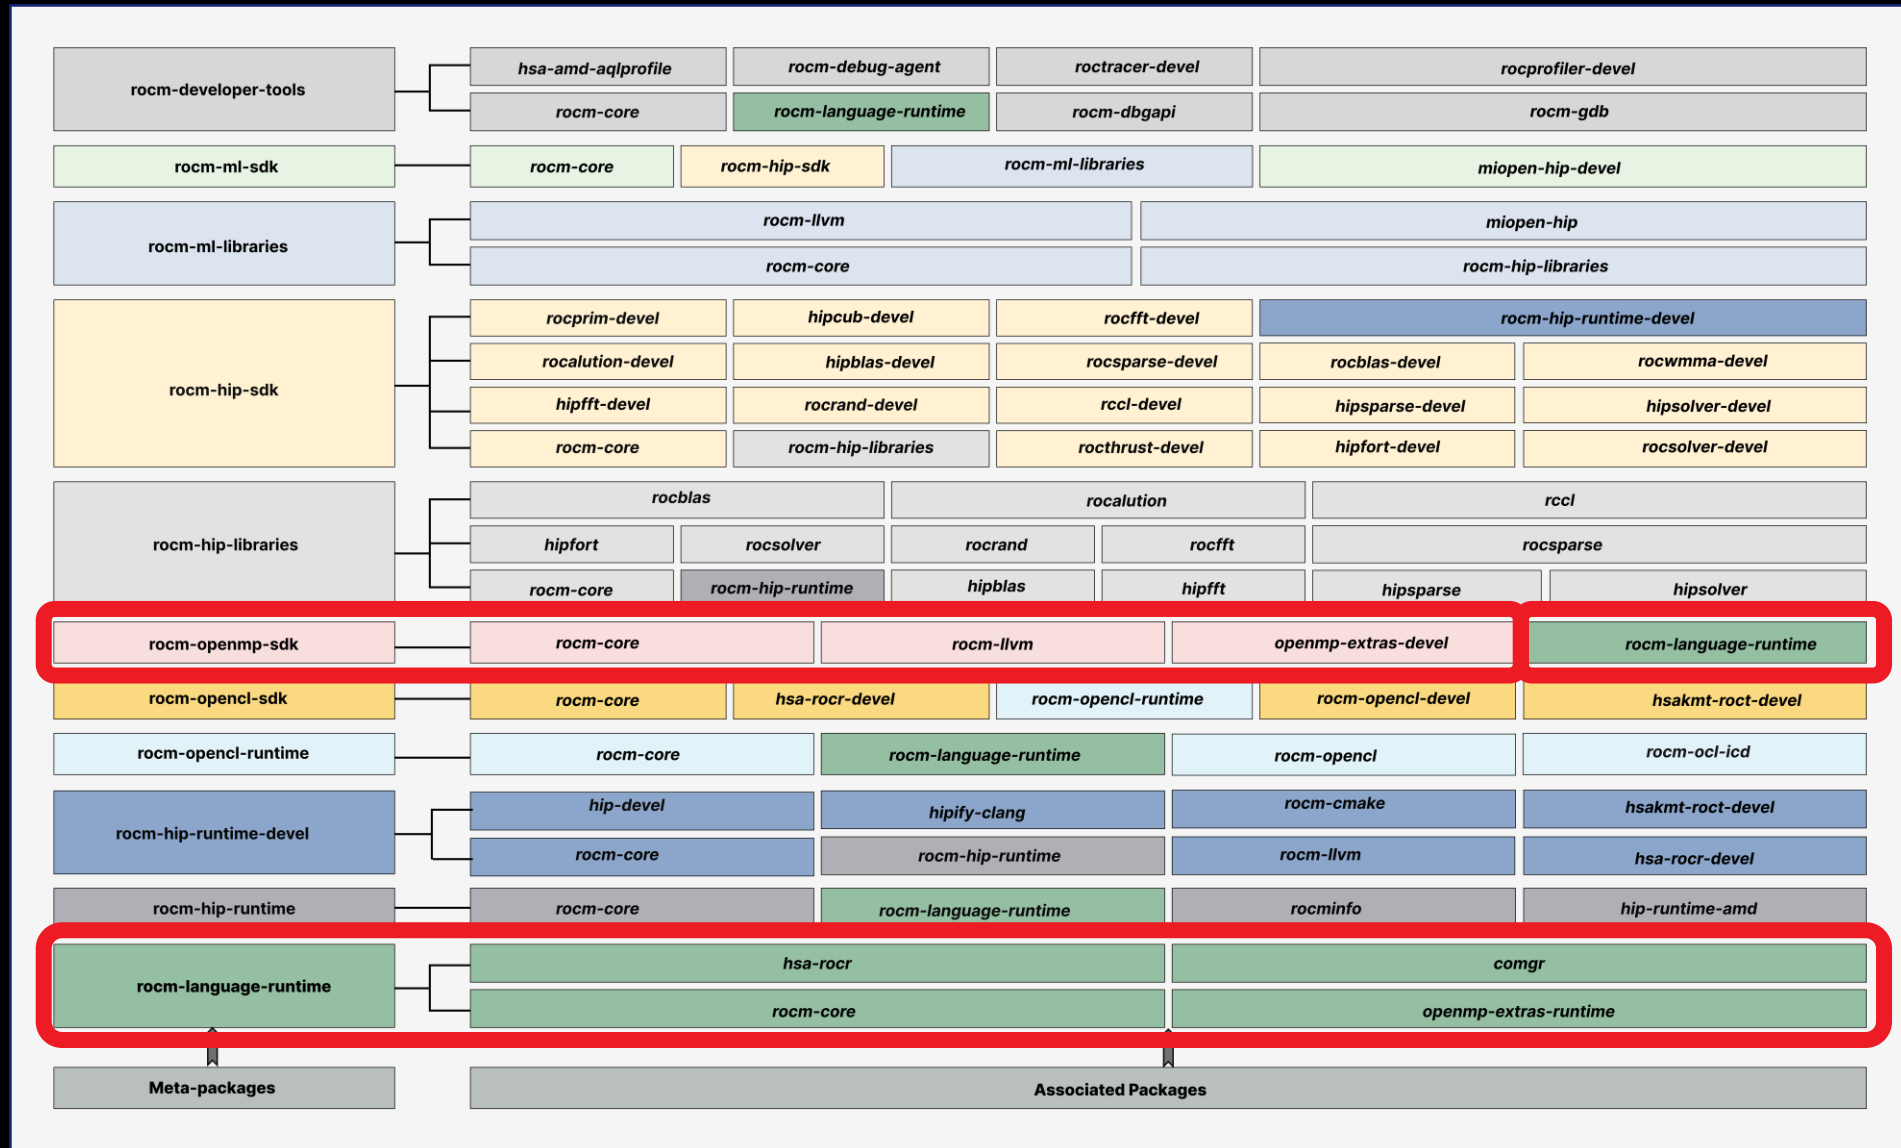
With AMD ROCm™ open software platform built for flexibility and performance, the HPC and AI communities can gain access to open compute languages, compilers, libraries and tools designed to accelerate code development and solve the toughest challenges in the world today.

# ROCm™ Software Stack (Meta Packages)



- ROCm ML Programming Model
- ROCm OpenMP Programming Model
- ROCm Opencil Programming Model
- ROCm HIP Programming Model
- ROCm llvm
- ROCm Components
- ROCm Components Categories

# ROCm™ Associated Packages



# Software System Terminology


## ROCm™



**AMD ROCm™ is an open software platform for GPU compute consisting of compilers, libraries, tools, etc.**

<https://www.amd.com/en/graphics/servers-solutions-rocm>


## AOMP



**AOMP is an open-source Clang/LLVM-based compiler with added support for the OpenMP® API on Radeon™ GPUs that builds on top of ROCm™**

Releases are more frequent than ROCm™  
(<https://github.com/ROCm-Developer-Tools/aomp>)

## LLVM



**Since AOMP is Clang/LLVM based, the high-level software architecture of OpenMP target offload support is identical**

AOMP is both ahead and behind LLVM mainline trunk (upstream)



# Jargon

- Upstream: LLVM mainline trunk (i.e., the code in their GitHub repository)
- Host: The host machine (e.g., AMD Epyc processor-based server)
- (Target) Device: An attached accelerator (e.g., AMD Instinct MI200)
- (Host) Runtime: The OpenMP runtime implementation for the host
- Libomptarget: The OpenMP target offload runtime (running on the host)
- (Target or Device) Runtime: The OpenMP runtime implementation for the target device
- Plugin: The device-specific target-runtime plugin
- Kernel: A target region executable on the device
  
- *Queue: An HSA (hardware) queue*
- *Signal: An HSA signal*

# AOMP

Open Source OpenMP Compiler  
for AMD GPUs



# AOMP

- **Open Source** OpenMP compiler for AMD GPUs
  - Available from <https://github.com/ROCm-Developer-Tools/aomp>
- Based on Clang/LLVM
  - AOMP tracks LLVM upstream closely, typically only a few hours behind
  - Includes additional optimizations (e.g., faster reductions)
  - Includes additional features (e.g., OMPT support)
  - Fixes are commonly submitted to upstream LLVM
- Fortran support
  - Currently via classic Flang
  - Active development on LLVM Flang



# AOMP

Getting, building and installing AOMP (and its dependencies)

- AOMP is “standalone” → requires only the kernel module (dkms) and libdrm
- AOMP is isolated from ROCm™ installations by installing in /usr/lib/aomp and uses RPATH on runtime libs
- AOMP includes builds of related ROCm™ components
  
- Packages are provided on Github for
  - CentOS 7, 8, and 9
  - SLES 15 SP 4
  - Ubuntu 20.04 and 22.04

# Components from ROCm™

- AOMP includes various ROCm™ components
  - The components are built from source
  - Versions and repositories are encoded in manifest file in the Github repository
  - Build scripts for individual components build the dependencies
  - All components are installed in AOMP install directory

## ROCm™ Components

- rocm-compiler-support
- rocm-device-libs
- ROCprofiler
- ROCtracer
- ROCdbgapi
- ROCgdb
- Hipamd
- Hip
- ROCclr
- ROCm-OpenCL-runtime
- ROCminfo
- ROCm-cmake
- ROCR-Runtime
- ROCT-Thunk-Interface
- hipfort

# AOMP Manifest File

```

Code Blame 33 lines (28 loc) · 2.94 KB
Raw Copy Download Edit View

1 <?xml version="1.0" encoding="UTF-8"?>
2 <manifest>
3   <!-- Manifest for AOMP 17.0-0 which uses ROCm 5.4 release branches of external repositories -->
4
5   <remote name="gerritgit" review="git.amd.com:8080" fetch="ssh://gerritgit/" />
6   <default revision="release/rocm-rel-5.4" remote="gerritgit" sync-j="4" sync-c="true" />
7   <remote name="roctools" fetch="https://github.com/ROCm-Developer-Tools/" />
8   <remote name="roc" fetch="https://github.com/RadeonOpenCompute/" />
9   <remote name="rocswh" fetch="https://github.com/ROCmSoftwarePlatform/" />
10
11  <project remote="roc" path="llvm-project" name="llvm-project" revision="aomp-17.0-0" groups="unlocked" />
12
13  <project remote="roc" path="rocm-compilersupport" name="ROCm-CompilerSupport" upstream="amd-stg-open" revision="38d950e7e7
14  <project remote="roc" path="rocm-device-libs" name="ROCm-Device-Libs" upstream="amd-stg-open" revision="07b347366e
15
16  <project remote="roctools" path="flang" name="flang" revision="aomp-17.0-0" groups="unlocked" />
17  <project remote="roctools" path="aomp-extras" name="aomp-extras" revision="aomp-17.0-0" groups="unlocked" />
18  <project remote="roctools" path="aomp" name="aomp" revision="aomp-17.0-0" groups="unlocked" />
19
20  <project remote="roctools" path="rocprofiler" name="rocprofiler" revision="rocm-5.4.3" groups="unlocked" />
21  <project remote="roctools" path="roctracer" name="roctracer" revision="rocm-5.4.3" groups="unlocked" />
22  <project remote="roctools" path="ROCdbgapi" name="ROCdbgapi" revision="rocm-5.4.3" groups="unlocked" />
23  <project remote="roctools" path="ROCgdb" name="ROCgdb" revision="rocm-5.4.3" groups="unlocked" />
24  <project remote="roctools" path="hipamd" name="hipamd" revision="rocm-5.4.3" groups="unlocked" />
25  <project remote="roctools" path="hip" name="hip" revision="rocm-5.4.3" groups="unlocked" />
26  <project remote="roctools" path="ROCclr" name="ROCclr" revision="rocm-5.4.3" groups="unlocked" />
27  <project remote="roc" path="ROCm-OpenCL-Runtime" name="ROCm-OpenCL-Runtime" revision="rocm-5.4.3" groups="unlocked" />
28  <project remote="roc" path="rocm-info" name="rocm-info" revision="rocm-5.4.3" groups="unlocked" />
29  <project remote="roc" path="rocm-cmake" name="rocm-cmake" revision="rocm-5.4.3" groups="unlocked" />

```

Used to tag dependencies for each AOMP release

- Used by AOMP team for development on most recent development branches
- Specific versions for each public release

# AOMP

Compilation Process

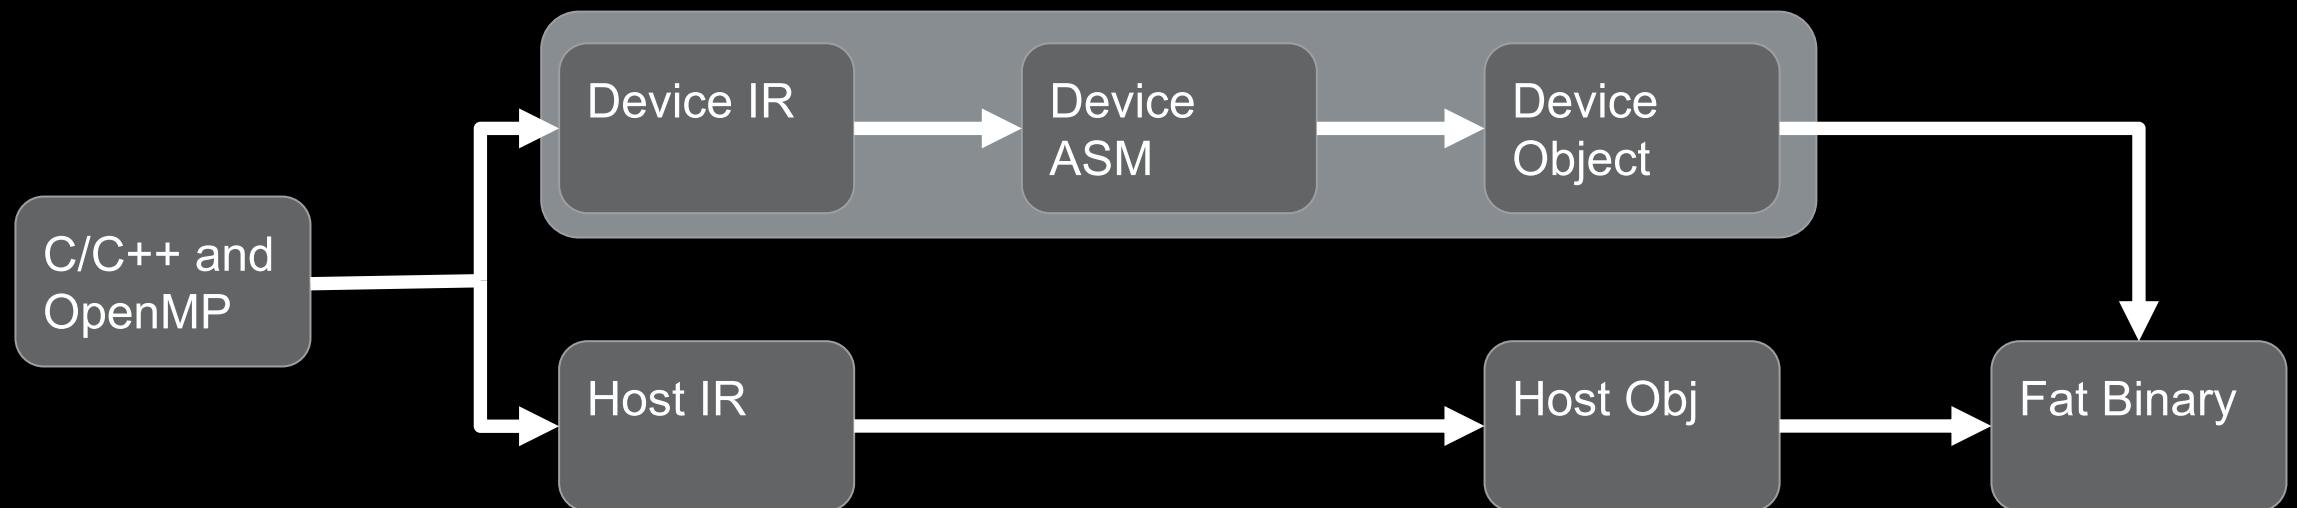
Execution

Device Plugin Architecture



# AOMP Compilation Process

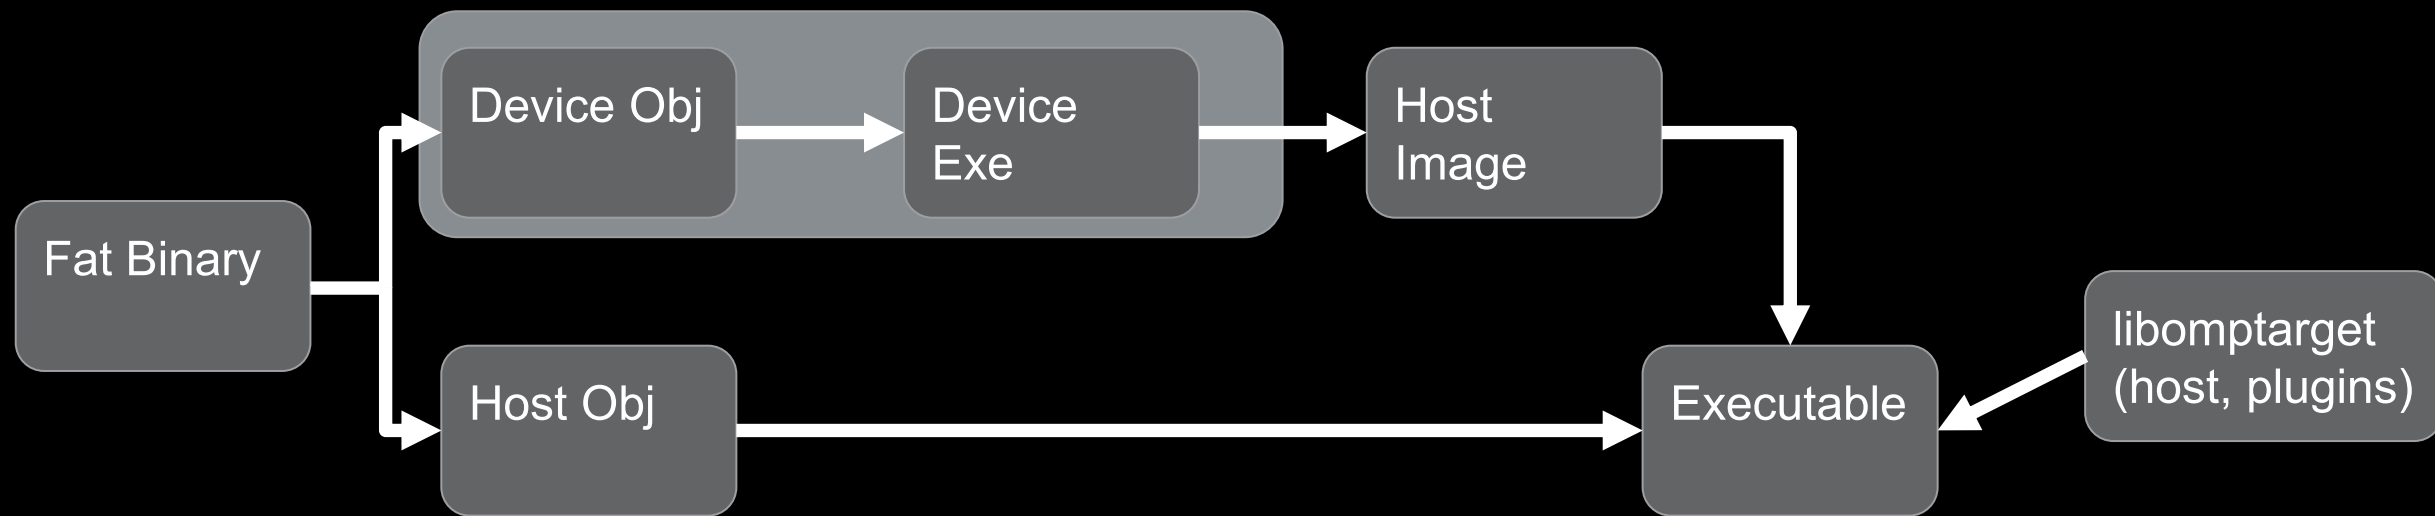
- Compilation requires code generation for host and for device
  - Invoke compiler twice for the different target triples
  - Object code for both targets is output into a single fat binary



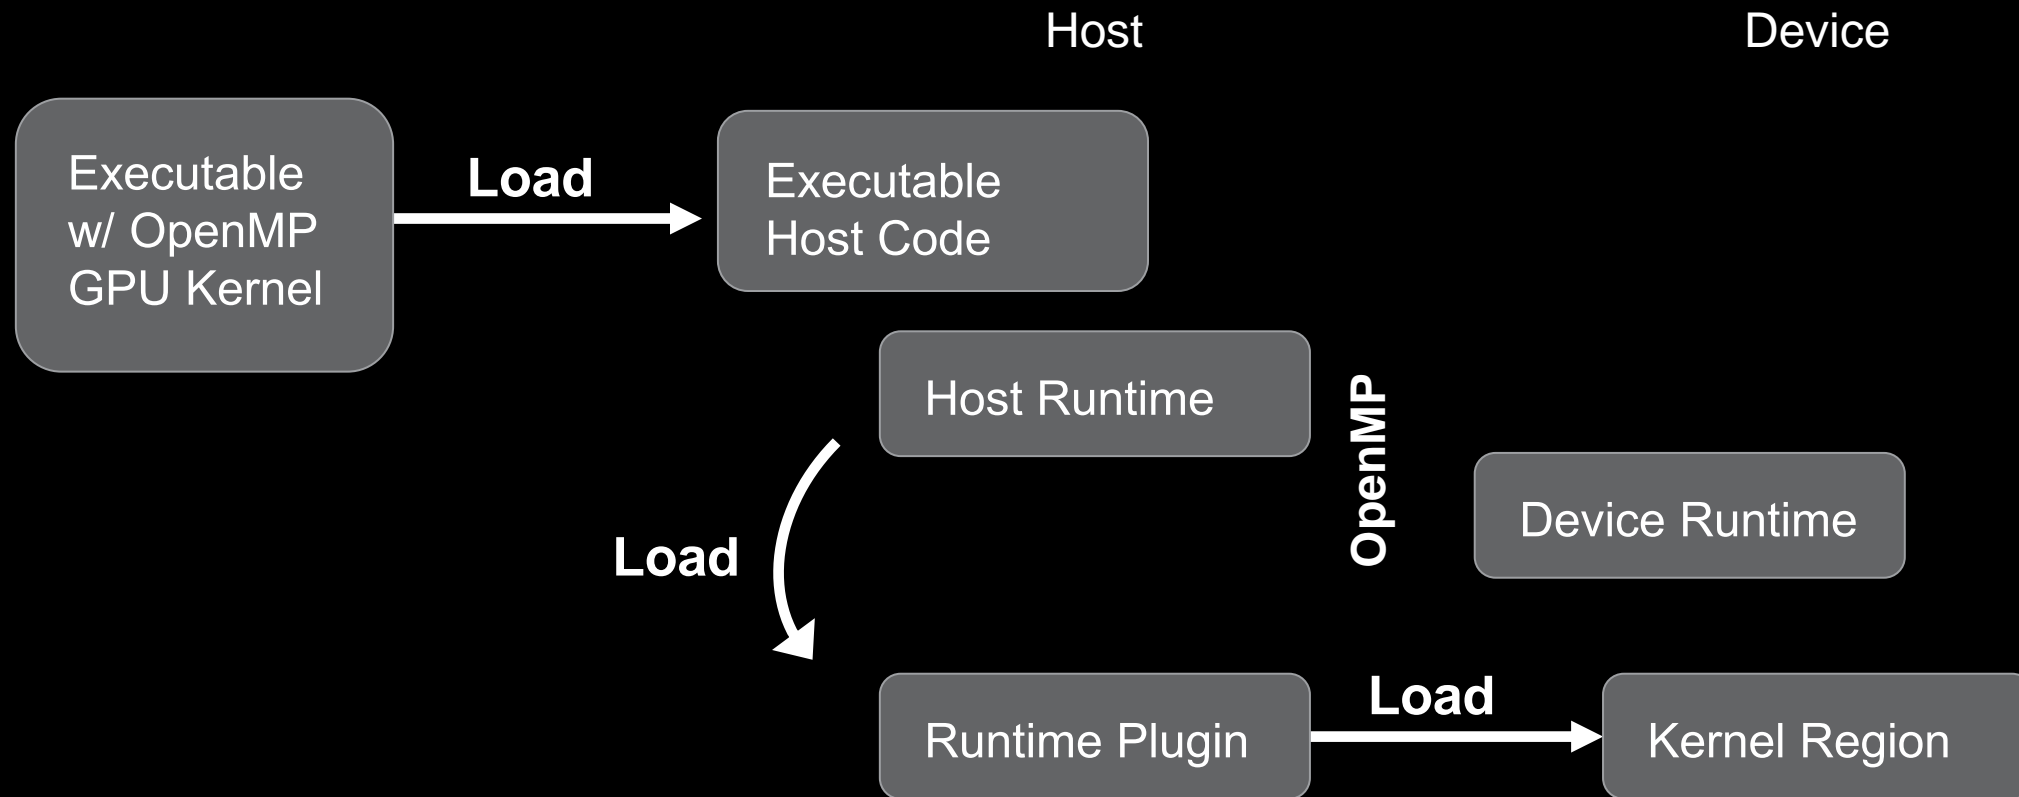


# AOMP Linking Process

- Linking requires generating one host image and an embedded device executable
  - Start with fat binary and *unbundle* into device and host objects again
  - Create a correctly linked host executable with embedded device ELF objects for offload kernels
- Moving towards a simplified toolchain, closer to today's upstream LLVM

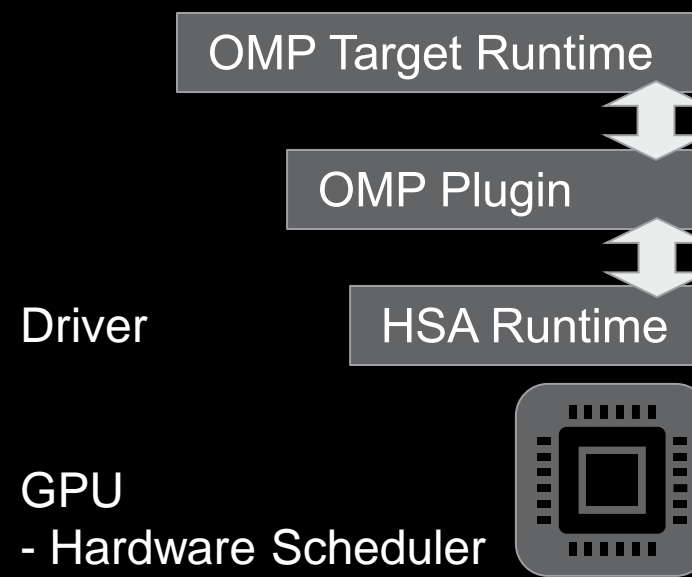


# High-level Runtime Process Overview



# AOMP Plugin Implementation using HSA

- HSA is the Heterogenous System Architecture, managed by the HSA foundation
  - Provides standardized interface to interact with heterogenous components
  - AMD provides vendor-specific extensions to the HSA standard as part of the ROCm™ stack
- AMDGPU Plugin builds on top of HSA making use of
  - HSA locked or pinned memory (i.e., non-migratable)
  - HSA Signals and Queues for (asynchronous) operation dispatch
  - HSA extensions to obtain profiling information for HSA signals
- HSA is low level!
  - `hsa_signal_create( ... )`
  - `hsa_signal_wait_relaxed( ... )`
  - `hsa_queue_load_write_index_acquire( ... )`
  - `hsa_queue_store_write_index_relaxed( ... )`



# Example of OpenMP Target Offload

# OpenMP Target Offload Code Example

```
int main(int argc, char **argv) {
    int vals[1024] = {0};

    #pragma omp target teams distribute parallel for map(vals)
    for(int i = 0; i < 1024; ++i) {
        vals[i] = 1;
    }

    for(const auto vi : vals) {
        std::cout << vi << '\n';
    }
    return 0;
}
```

# OpenMP Target Offload Code Example

```
int main(int argc, char **argv) {
    int vals[1024] = {0};

    #pragma omp target teams distribute parallel for map(tofrom: vals)
    for(int i = 0; i < 1024; ++i) {
        vals[i] = 1;
    }

    for(const auto vi : vals) {
        std::cout << vi << '\n';
    }
    return 0;
}
```

# OpenMP Target Offload Code Example

```
int main(int argc, char **argv) {  
    int vals[1024] = {0};
```

```
#pragma omp target teams distribute parallel for map(vals)  
for(int i = 0; i < 1024; ++i) {  
    vals[i] = 1;  
}
```

**Executes on  
the device**

```
for(const auto vi : vals) {  
    std::cout << vi << '\n';  
}  
return 0;  
}
```

# Libomptarget Plugin Entry Points

```
int main(int argc, char **argv) {  
    int vals[1024] = {0};
```

```
#pragma omp target teams distribute parallel for map(vals)
```

```
for(int i = 0; i < 1024; ++i) {  
    vals[i] = 1;  
}
```

```
for(const auto vi : vals) {  
    std::cout << vi << '\n';  
}  
return 0;
```

```
}
```

Data submit

Run region

Data retrieve



# AOMP Summary

- Open-source compiler for OpenMP target offload based on Clang/LLVM
  - Support for C/C++/Fortran
  - Features are merged into ROCm™ releases
- Additional optimizations compared to upstream Clang/LLVM
- Provides standalone build and installation from source using provided scripts
  - Requires the kernel driver to be installed on the system
  - Pulls-in the required ROCm™ dependencies

**AMD** 