

# AMD ROCm™ Basics & Optimization Overview

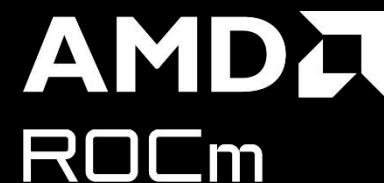
Joe Liu 刘仕洲  
Jan 2026

**AMD** together we advance\_

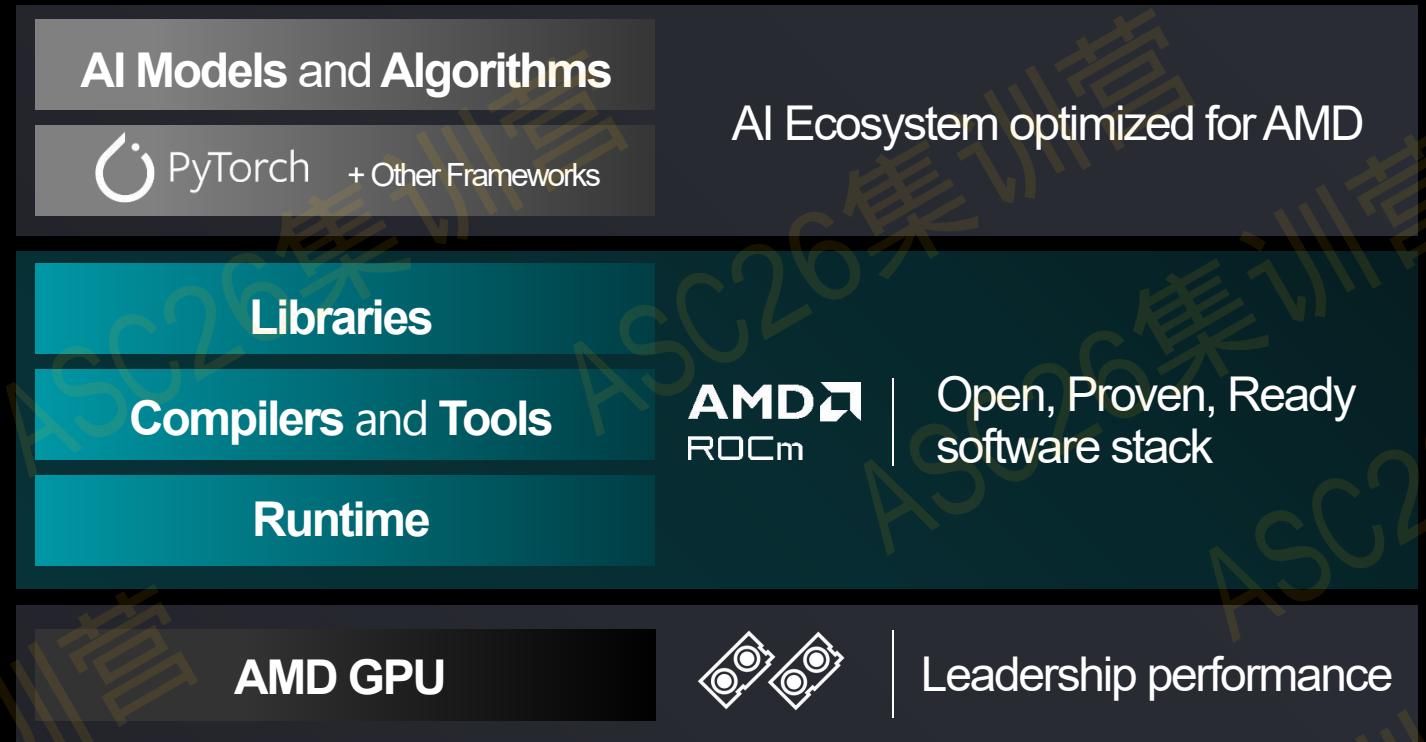
# Agenda

---

1. Introduction to the AMD ROCm™ Software Stack
2. Transitioning Workloads to AMD GPUs
3. Performance Optimization
  - Optimizing application using popular libraries
  - Profiling the models
  - Adding HIP kernel to implement a custom layer
4. Available Collaterals, Q&A



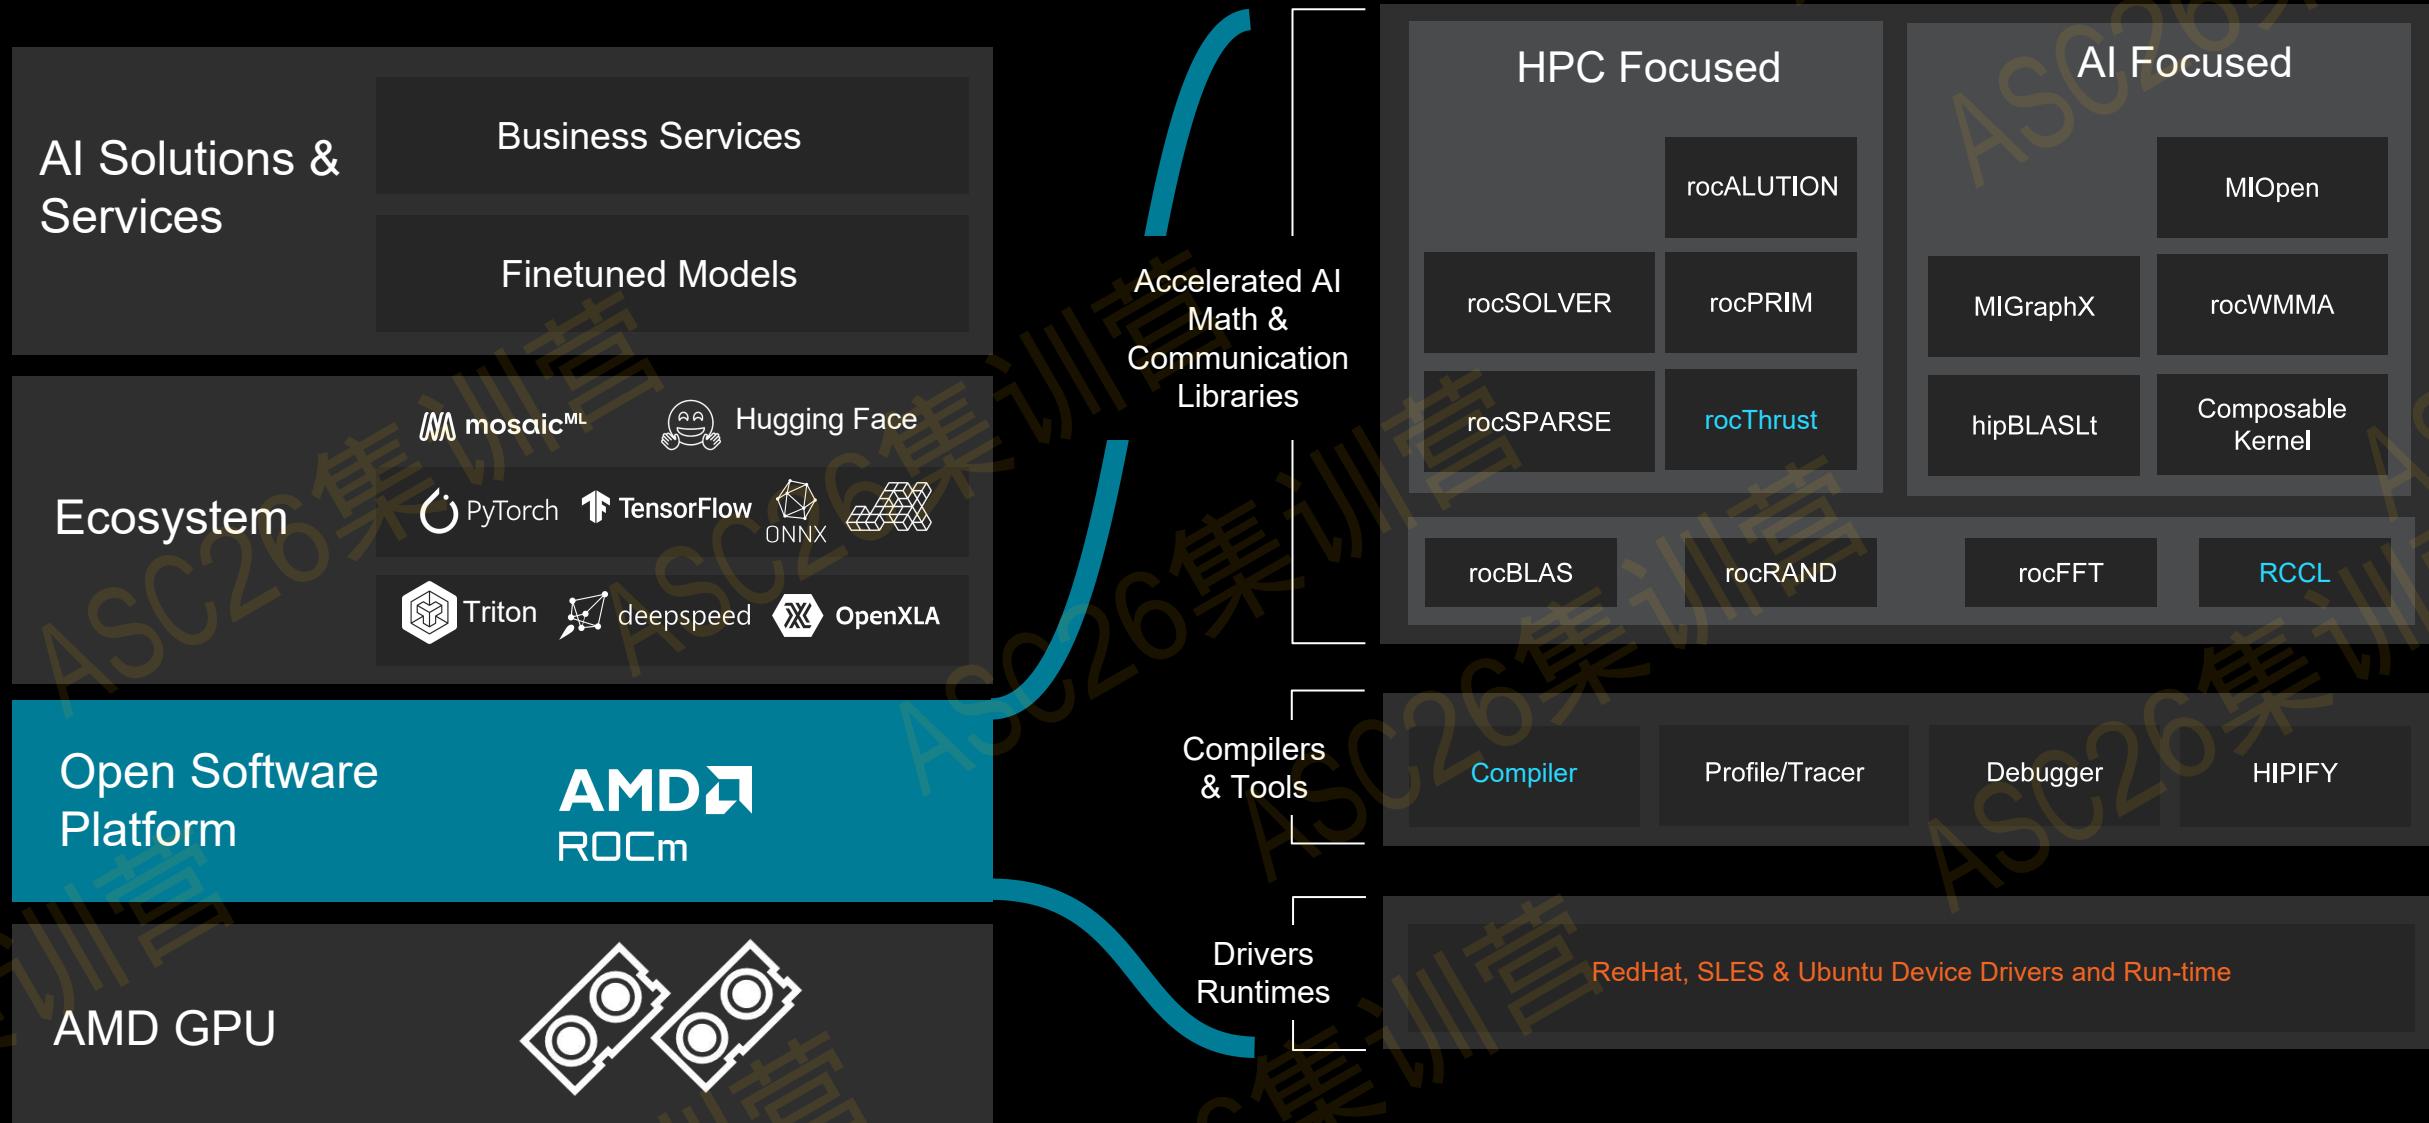
# Optimized AI Software Stack



- Commitment to **Open-Source**
- **No Code Change** Execution
- Optimized for **Generative AI**

# AMD ROCm™ Software Stack

MIT/BSD License  
Apache License  
GPL License

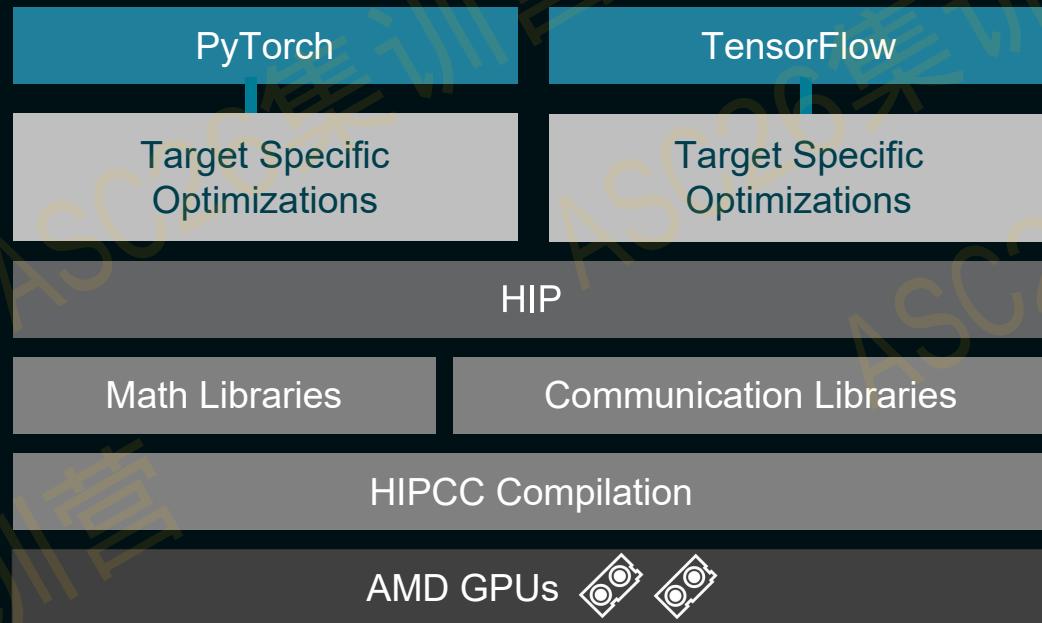


Use of third party marks/logos/products is for informational purposes only and no endorsement of or by AMD is intended or implied GD-83

# Library and Compiler Based Optimization

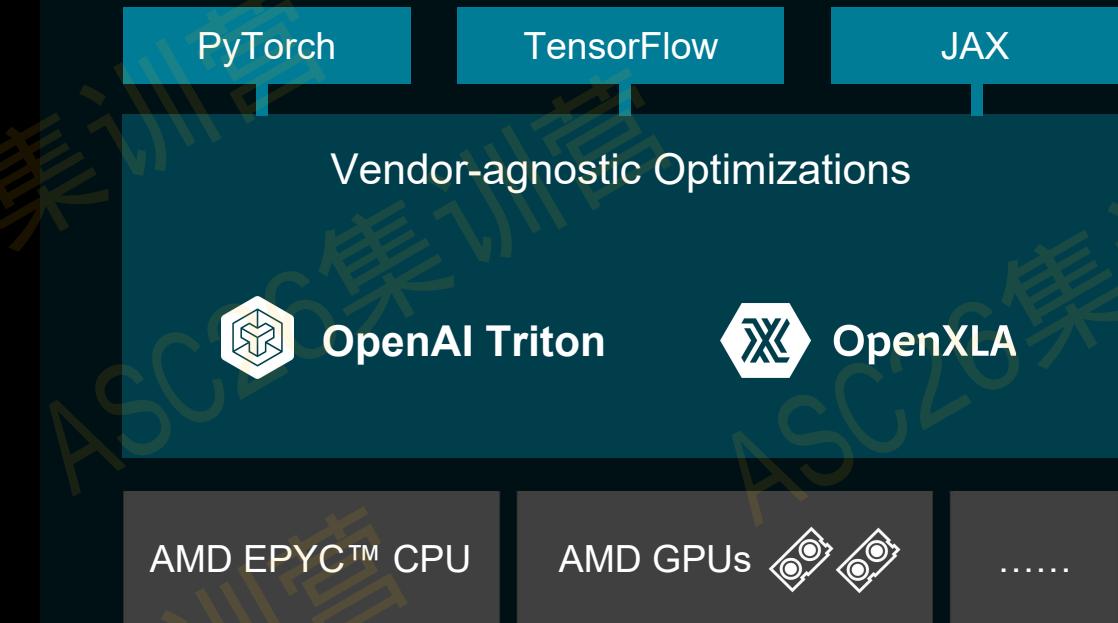
1

## Max Performance Framework Operator Optimization



2

## Max Portability IR-based Optimization

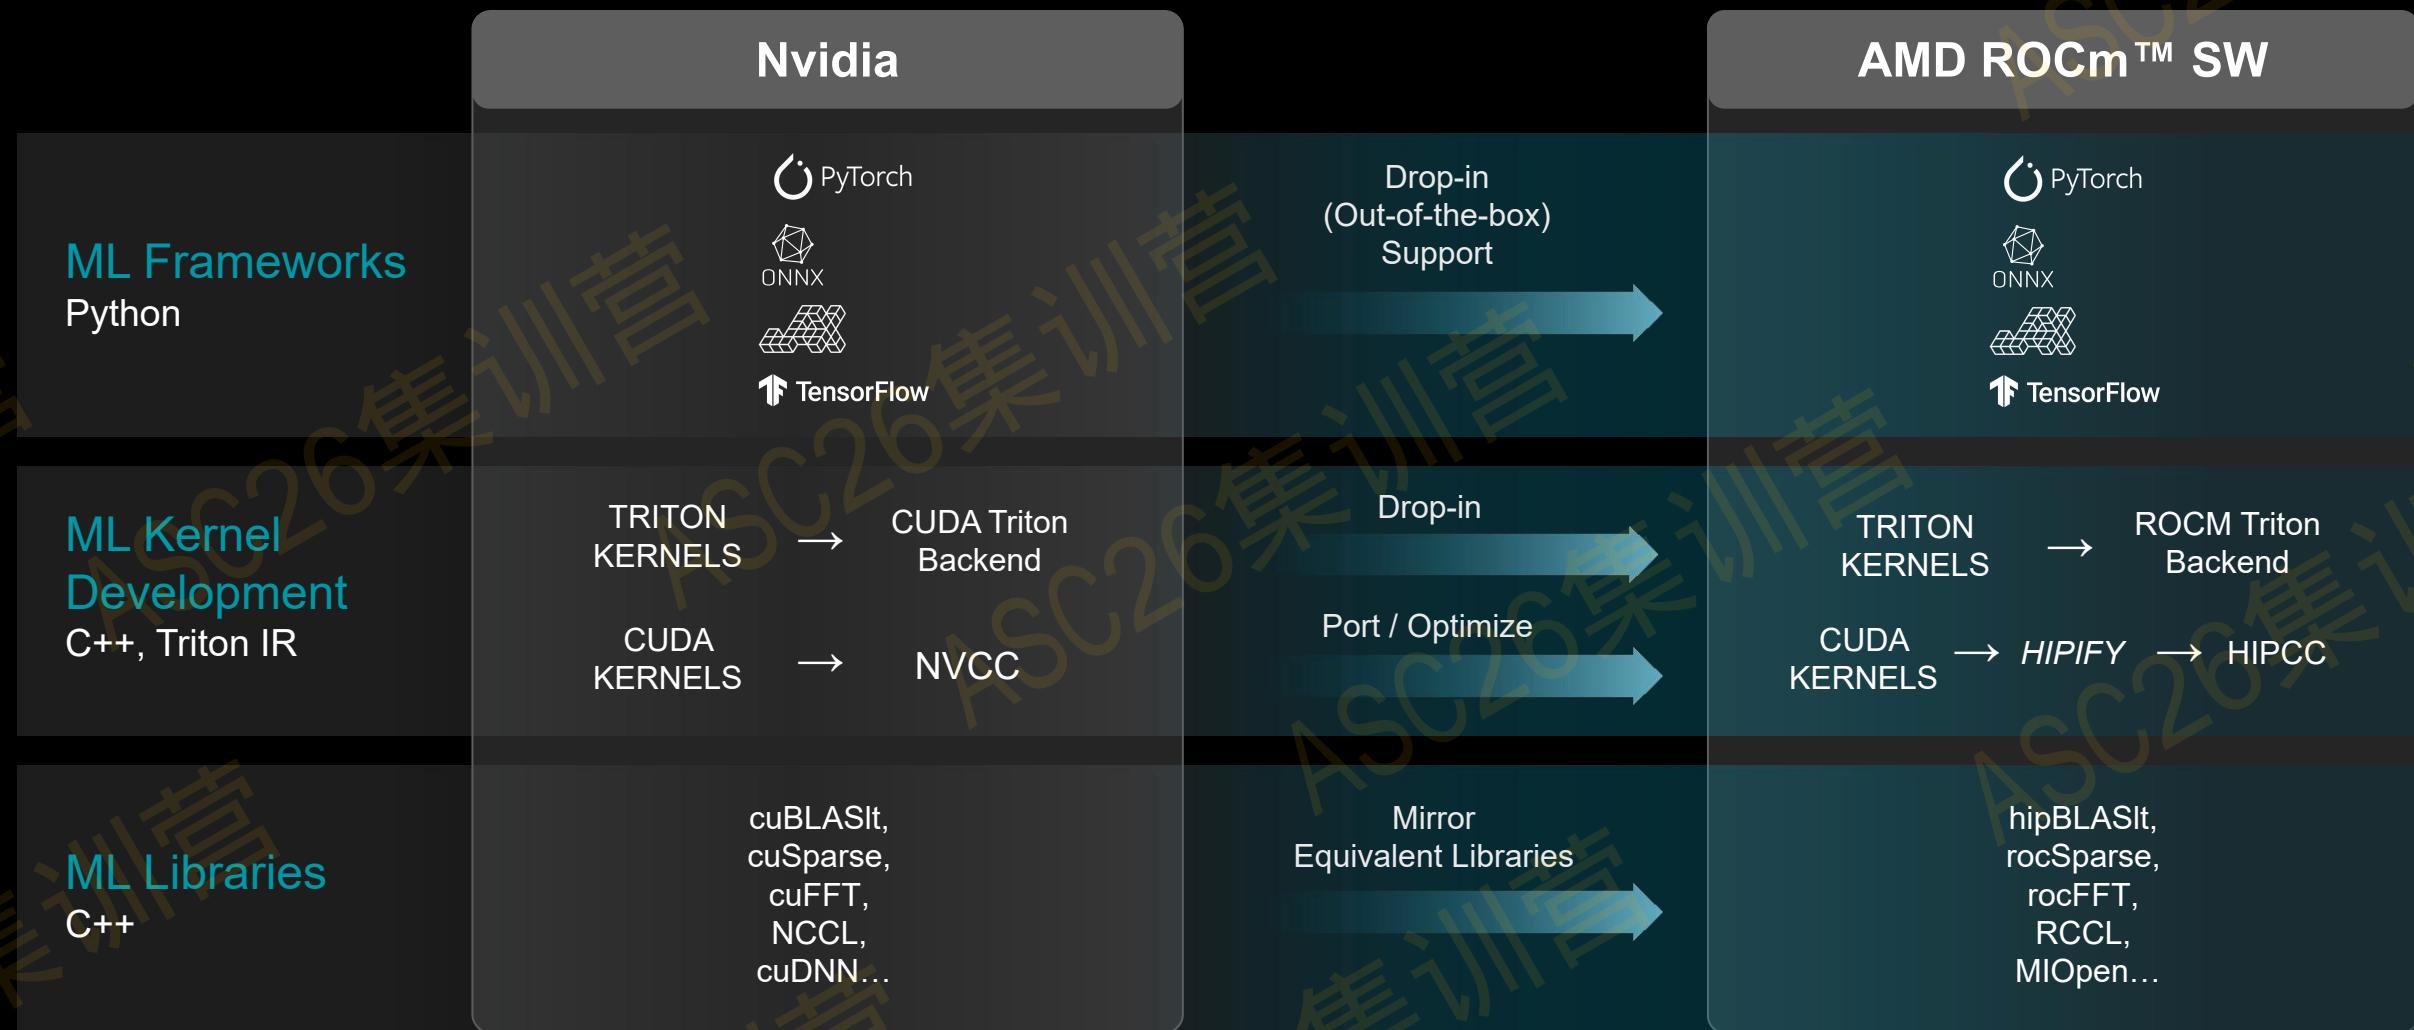


# Agenda

---

1. Introduction to the AMD ROCm™ Software Stack
2. **Transitioning Workloads to AMD GPUs**
3. Performance Optimization
  - Optimizing application using popular libraries
  - Profiling the models
  - Adding HIP kernel to implement a custom layer
4. Available Collaterals, Q&A

# Transitioning AI Workloads to AMD GPUs



# ROCM™ Software: Can You Spot a Difference?

NVIDIA CUDA

```
import torch
import torch.nn as nn

# Get cpu or gpu device for training.
device = "cuda:0" if torch.cuda.is_available() else "cpu"
print(f"Using {device} device")

# Define model
class Network(nn.Module):
    def __init__(self):
        super().__init__()
        self.flatten = nn.Flatten()
        self.linear_relu_stack = nn.Sequential(
            nn.Linear(28 * 28, 512),
            nn.ReLU(),
            nn.Linear(512, 512),
            nn.ReLU(),
            nn.Linear(512, 10)
        )

    def forward(self, x):
        x = self.flatten(x)
        logits = self.linear_relu_stack(x)
        return logits

model = Network().to(device)
print(model)
```



AMD ROCm™ Software

```
import torch
import torch.nn as nn

# Get cpu or gpu device for training.
device = "cuda:0" if torch.cuda.is_available() else "cpu"
print(f"Using {device} device")

# Define model
class Network(nn.Module):
    def __init__(self):
        super().__init__()
        self.flatten = nn.Flatten()
        self.linear_relu_stack = nn.Sequential(
            nn.Linear(28 * 28, 512),
            nn.ReLU(),
            nn.Linear(512, 512),
            nn.ReLU(),
            nn.Linear(512, 10)
        )

    def forward(self, x):
        x = self.flatten(x)
        logits = self.linear_relu_stack(x)
        return logits

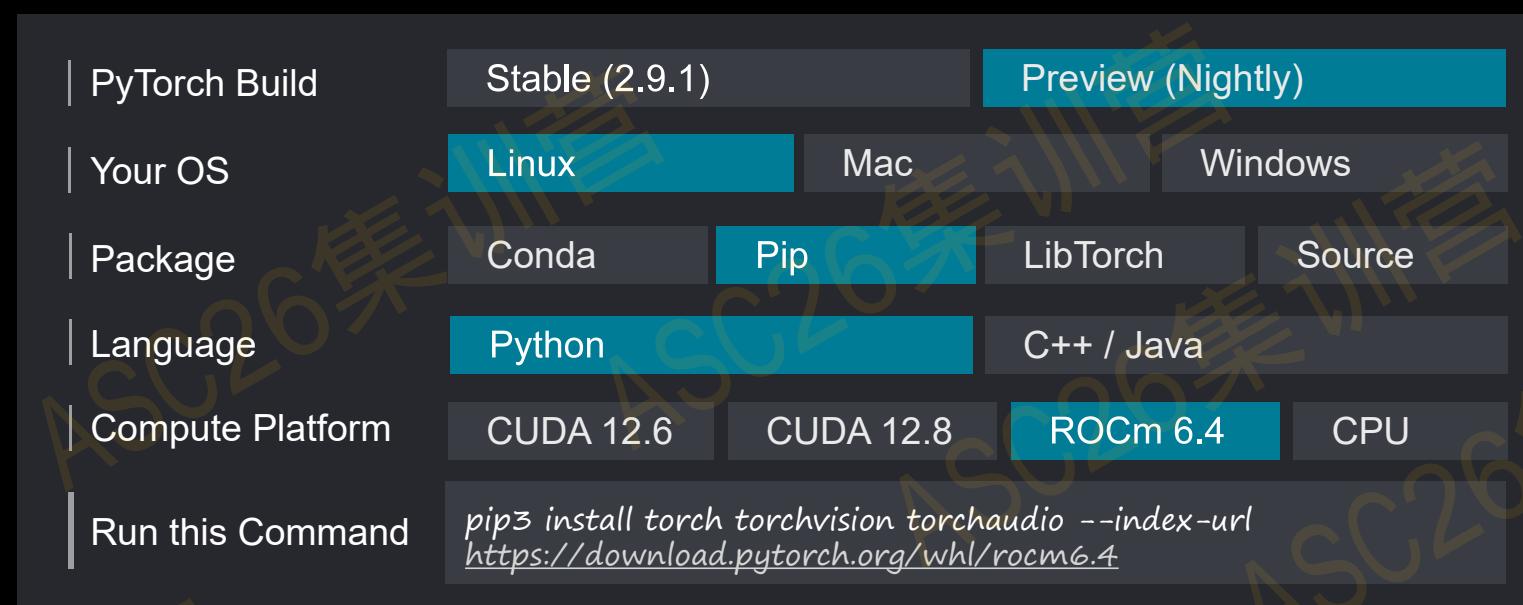
model = Network().to(device)
print(model)
```

# PyTorch 2.8 Easily Enabled on AMD GPUs

Step 1. Install ROCm™ Software (Driver and SDK)

Step 2. Install Pip Wheel From Pytorch.Org

Step 3. Run Existing Code -- No Changes Required



Optionally Install Docker containers from:

- `rocm/pytorch:latest`
- `rocm/pytorch-nightly:latest`

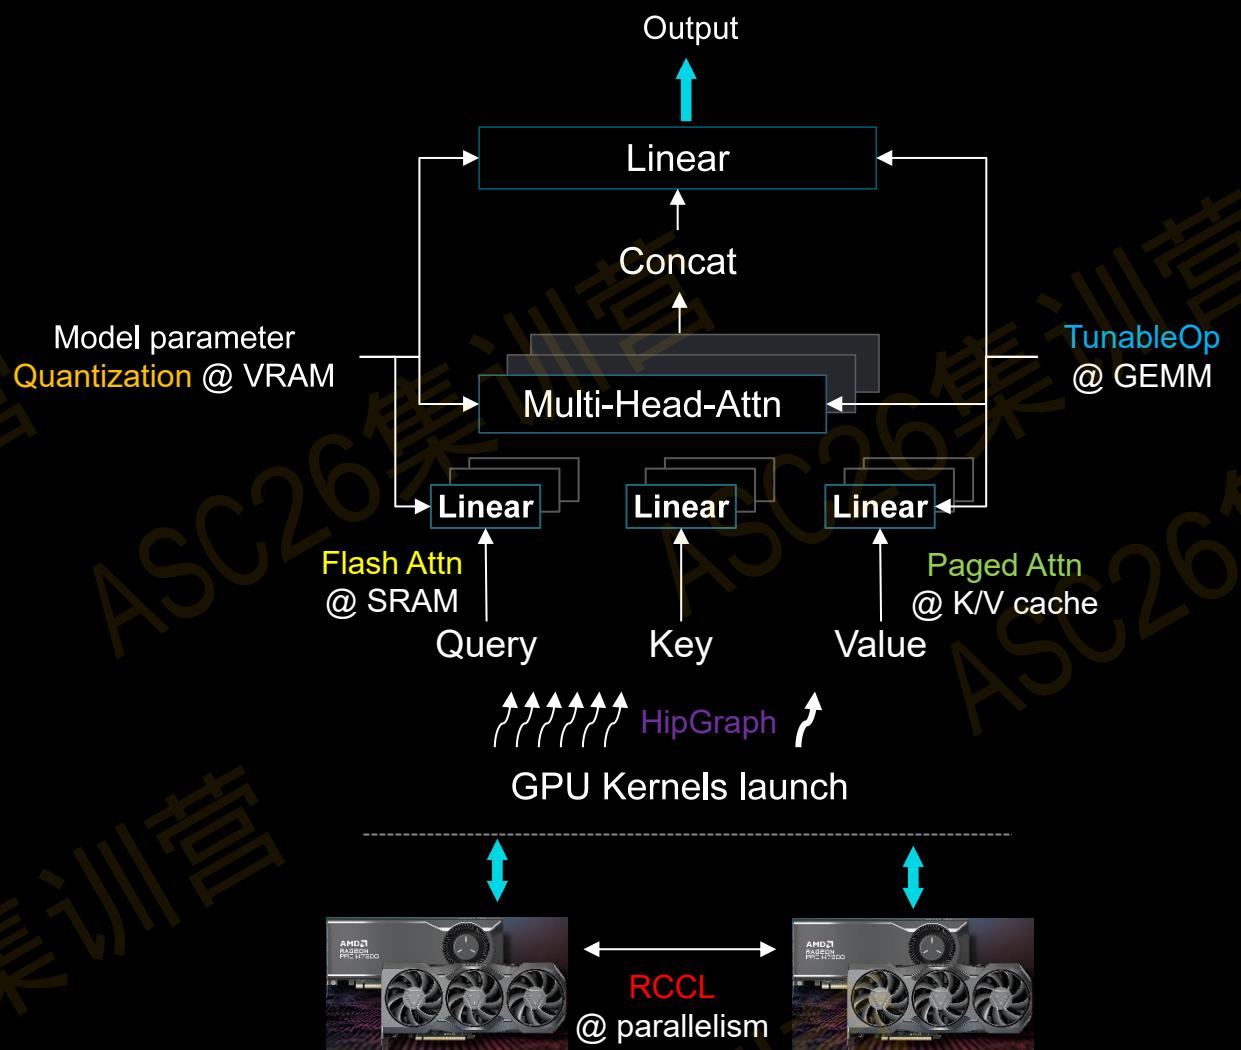
Use of third party marks/logos/products is for informational purposes only and no endorsement of or by AMD is intended or implied GD-83

# Agenda

---

1. Introduction to the AMD ROCm™ Software Stack
2. Transitioning Workloads to AMD GPUs
3. Performance Optimization
  - Optimizing application using popular libraries
  - Profiling the models
  - Adding HIP kernel to implement a custom layer
4. Available Collaterals, Q&A

# Inference Challenges and Optimization Opportunities



## Flash Attention, Xformers

- Tiling of input sequence in GPU SRAM to reduce VRAM data movement

## Paged Attention

- Partitioned KV cache into fixed size blocks to reduce memory usage

## GEMM Optimization – PyTorch TunableOp

- Automatic selection of the best performing GEMM kernels

## Graph Optimization – HipGraph

- Launch multiple kernels through a single CPU operation

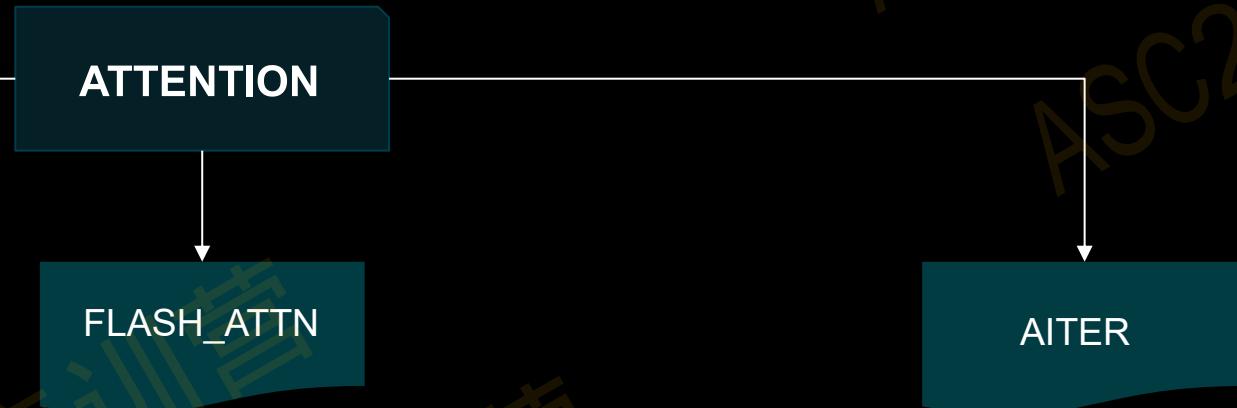
## Collective communication – RCCL

- Collective Ops across multiple devices to support Tensor/Pipeline parallel

## Quantization – GPTQ, Bitsandbytes

- Weight-only compression to reduce video memory footprint

# Portability - Libraries



```
import xformers.ops as xops

out = xops.memory_efficient_attention(q,
                                      k,
                                      v,
                                      attn_bias=None,
                                      op        =None)
```

```
from flash_attn import flash_attn_varlen_func
# batch and sequence dimensions merged into a single dimension
q, k, v = (rearrange(x, "b s ... -> (b s) ...")
            for x in [q, k, v])
out = flash_attn_varlen_func(q,
                             k,
                             v,
                             cu_seqlens_q=cu_seqlens,
                             cu_seqlens_k=cu_seqlens,
                             max_seqlen_q=max_seqlen,
                             max_seqlen_k=max_seqlen)
```

```
from aiter import flash_attn_varlen_func
# batch and sequence dimensions merged into a single dimension
q, k, v = (rearrange(x, "b s ... -> (b s) ...")
            for x in [q, k, v])
out = flash_attn_varlen_func(q,
                             k,
                             v,
                             cu_seqlens_q=cu_seqlens,
                             cu_seqlens_k=cu_seqlens,
                             max_seqlen_q=max_seqlen,
                             max_seqlen_k=max_seqlen)
```

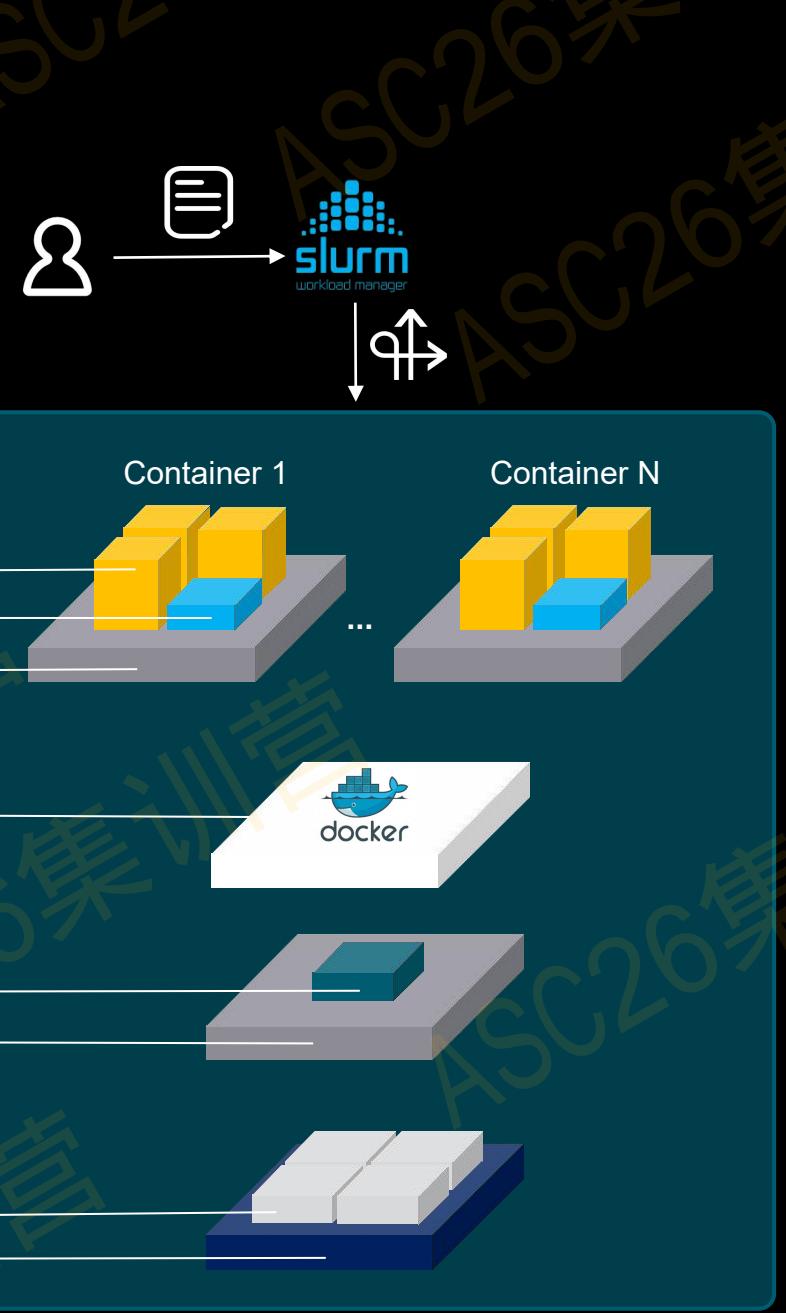
# Agenda

---

1. Introduction to the AMD ROCm™ Software Stack
2. Transitioning Workloads to AMD GPUs
3. Performance Optimization
  - Optimizing application using popular libraries
  - [Profiling the models](#)
  - Adding HIP kernel to implement a custom layer
4. Available Collaterals, Q&A

# The Components in the Environment

- User submits jobs (sbatch / srun)
- Slurm scheduling layer
  - Allocate nodes / CPU / GPU
  - Launch the container runtime
- Container layer (Docker / Apptainer)
  - Application
  - ROCm user space (HIP Runtime / rocBLAS / MIOpen)
- Host driver layer
  - ROCm driver + kernel
  - /dev/kfd, /dev/dri device mapping
- Hardware layer
  - AMD GPUs & CPUs



# The Profiling Tools and Visualization - rocm-smi

- A command-line utility and library provided by ROCm for monitoring the following AMD GPU status:
  - Power, temperature, clocks (gfx/mem), voltage, fan speed
  - GPU utilization, memory usage (VRAM/GTT), PCIe link speed/width

```

# Show a quick summary of all GPUs
rocm-smi

# Detailed power, temps, clocks, and utilization
rocm-smi --showpower --showtemp --showclocks --showuse

# Memory usage and PCIe info
rocm-smi --showmemuse -showbus

# List GPU processes
rocm-smi -showpids

# Real-time monitoring (refresh every 0.1s)
watch -n 0 rocm-smi

```

**watch -c rocm-smi --showclocks**

Every 2.0s: rocm-smi --showclocks

```

=====
ROCM System Management Interface =====
===== Current clock frequencies =====
GPU[0] : dcefclk clock level: 0: (145MHz)
GPU[0] : fclk clock level: 1: (1000MHz)
GPU[0] : mclk clock level: 0: (96MHz)
GPU[0] : sclk clock level: 1: (0MHz)
GPU[0] : socclk clock level: 1: (600MHz)
GPU[0] : pcie clock level: 0 (16.0GT/s x16)
GPU[1] : dcefclk clock level: 0: (145MHz)
GPU[1] : fclk clock level: 1: (1000MHz)
GPU[1] : mclk clock level: 0: (96MHz)
GPU[1] : sclk clock level: 1: (0MHz)
GPU[1] : socclk clock level: 1: (600MHz)
GPU[1] : pcie clock level: 0 (16.0GT/s x16)
GPU[2] : dcefclk clock level: 0: (145MHz)
GPU[2] : fclk clock level: 1: (1000MHz)
GPU[2] : mclk clock level: 0: (96MHz)
GPU[2] : sclk clock level: 1: (0MHz)
GPU[2] : socclk clock level: 1: (600MHz)
GPU[2] : pcie clock level: 0 (16.0GT/s x16)
GPU[3] : dcefclk clock level: 0: (145MHz)
GPU[3] : fclk clock level: 1: (1000MHz)
GPU[3] : mclk clock level: 0: (96MHz)
GPU[3] : sclk clock level: 1: (0MHz)
GPU[3] : socclk clock level: 1: (600MHz)
GPU[3] : pcie clock level: 0 (16.0GT/s x16)
GPU[4] : dcefclk clock level: 0: (145MHz)
GPU[4] : fclk clock level: 1: (1000MHz)
GPU[4] : mclk clock level: 0: (96MHz)
GPU[4] : sclk clock level: 1: (0MHz)
GPU[4] : socclk clock level: 1: (600MHz)
GPU[4] : pcie clock level: 0 (16.0GT/s x16)
GPU[5] : dcefclk clock level: 0: (145MHz)
GPU[5] : fclk clock level: 1: (1000MHz)
GPU[5] : mclk clock level: 0: (96MHz)
GPU[5] : sclk clock level: 1: (0MHz)
GPU[5] : socclk clock level: 1: (600MHz)
GPU[5] : pcie clock level: 0 (16.0GT/s x16)
GPU[6] : dcefclk clock level: 0: (145MHz)
GPU[6] : fclk clock level: 1: (1000MHz)
GPU[6] : mclk clock level: 0: (96MHz)
GPU[6] : sclk clock level: 1: (0MHz)
GPU[6] : socclk clock level: 1: (600MHz)
GPU[6] : pcie clock level: 0 (16.0GT/s x16)
GPU[7] : dcefclk clock level: 0: (145MHz)
GPU[7] : fclk clock level: 0: (601MHz)
GPU[7] : mclk clock level: 0: (96MHz)
GPU[7] : sclk clock level: 1: (0MHz)
GPU[7] : socclk clock level: 0: (500MHz)
GPU[7] : pcie clock level: 0 (16.0GT/s x16)

=====
End of ROCm SMI Log =====

```

# The Profiling Tools and Visualization



PyTorch

## PyTorch Profiler

- [https://pytorch.org/tutorials/recipes/recipes/profiler\\_recipe.html](https://pytorch.org/tutorials/recipes/recipes/profiler_recipe.html)

```
import torch
from torch.profiler import profile, record_function, ProfilerActivity
```



## ROCKProfiler

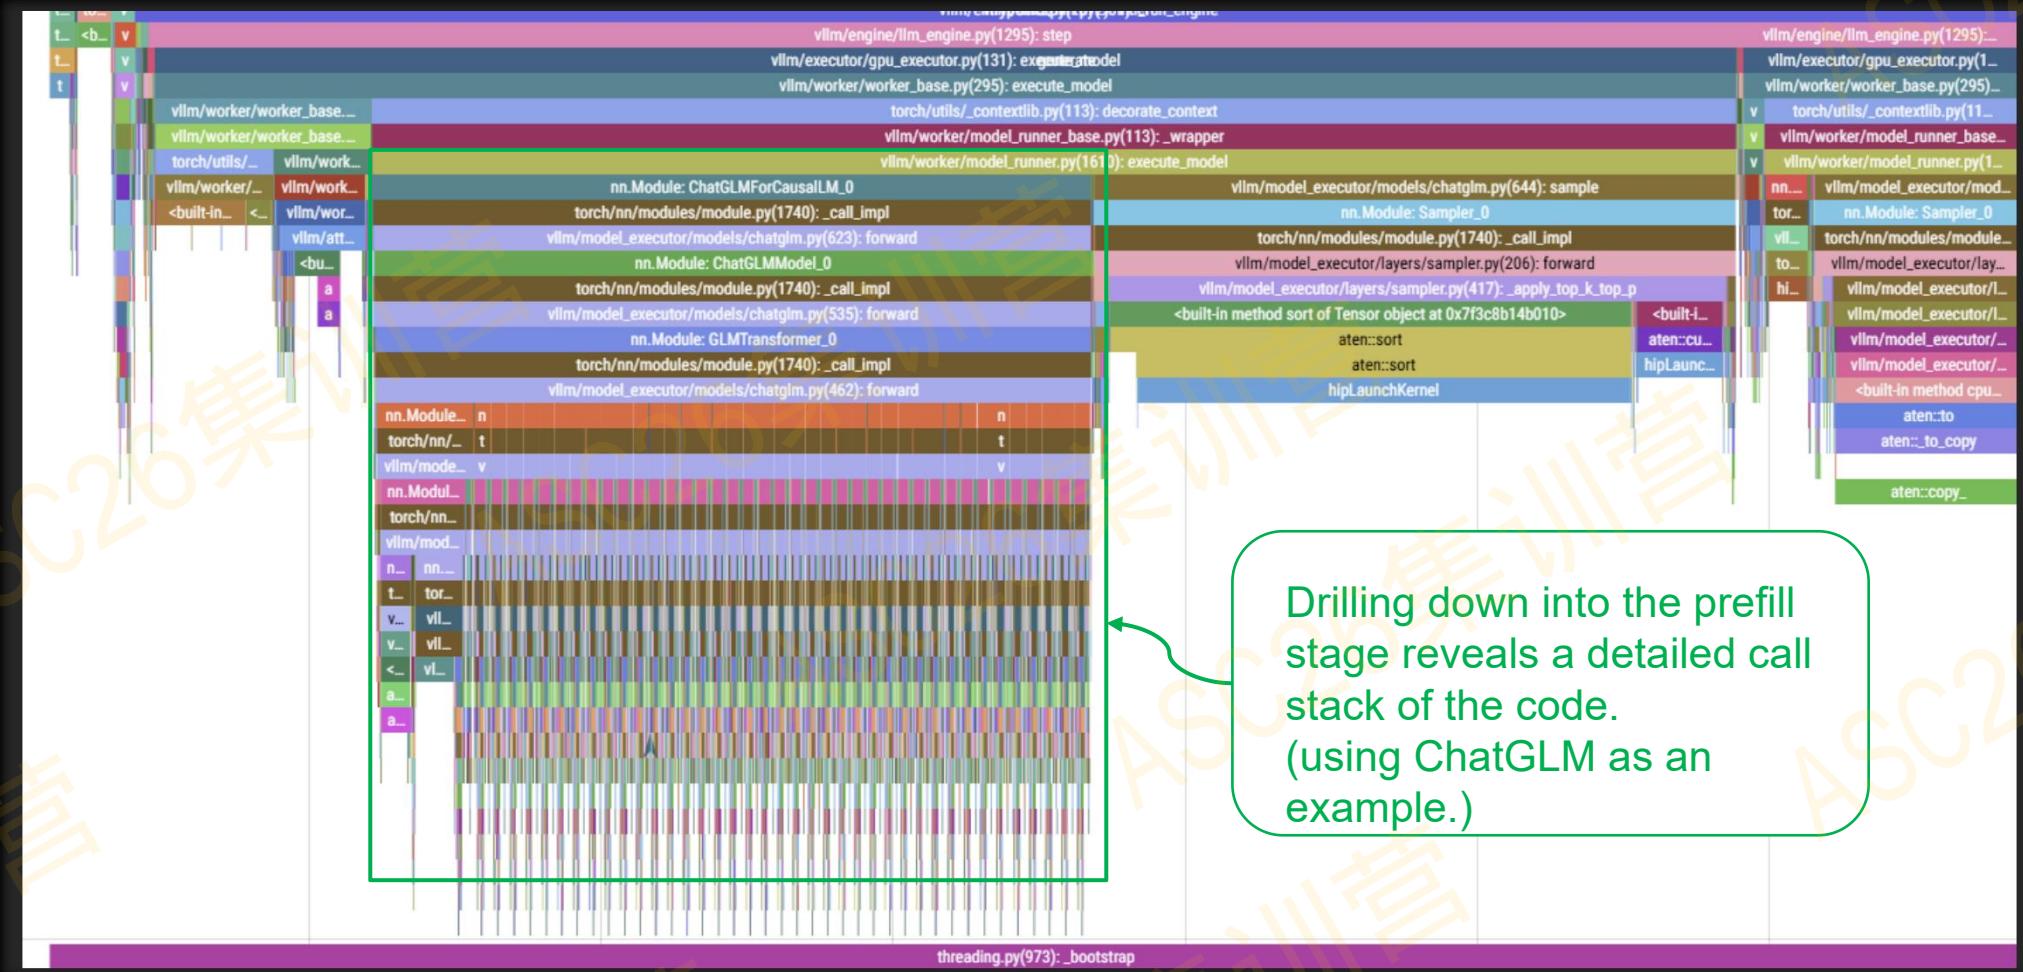
- <https://rocm.docs.amd.com/projects/rocprofiler/en/latest/install/install.html>
- *rocprof* and *rocprofv2* are included as standard components of the ROCm distribution

```
rocprof -d outputFolder --hip-trace ./Matrixtranspose
```

- *ROCTracer API* is a library that requires minor code modification in the application to be traced but provides greater flexibility

Use of third party marks/logos/products is for informational purposes only and no endorsement of or by AMD is intended or implied GD-83

# The Profiling Tools and Visualization - Samples



Input Plots Options

Counters Wave States Hotspot Occupancy Kernel Dispatches Compute Unit Utilization

UI Path:  
tput\_agent\_50641\_dispatch\_200033Shader SIMD Slot WaveID  
0 0 0 0

GlobalView zoom: 10

WaveView zoom: 10

Iteration 1

WaveView clock range:

33172

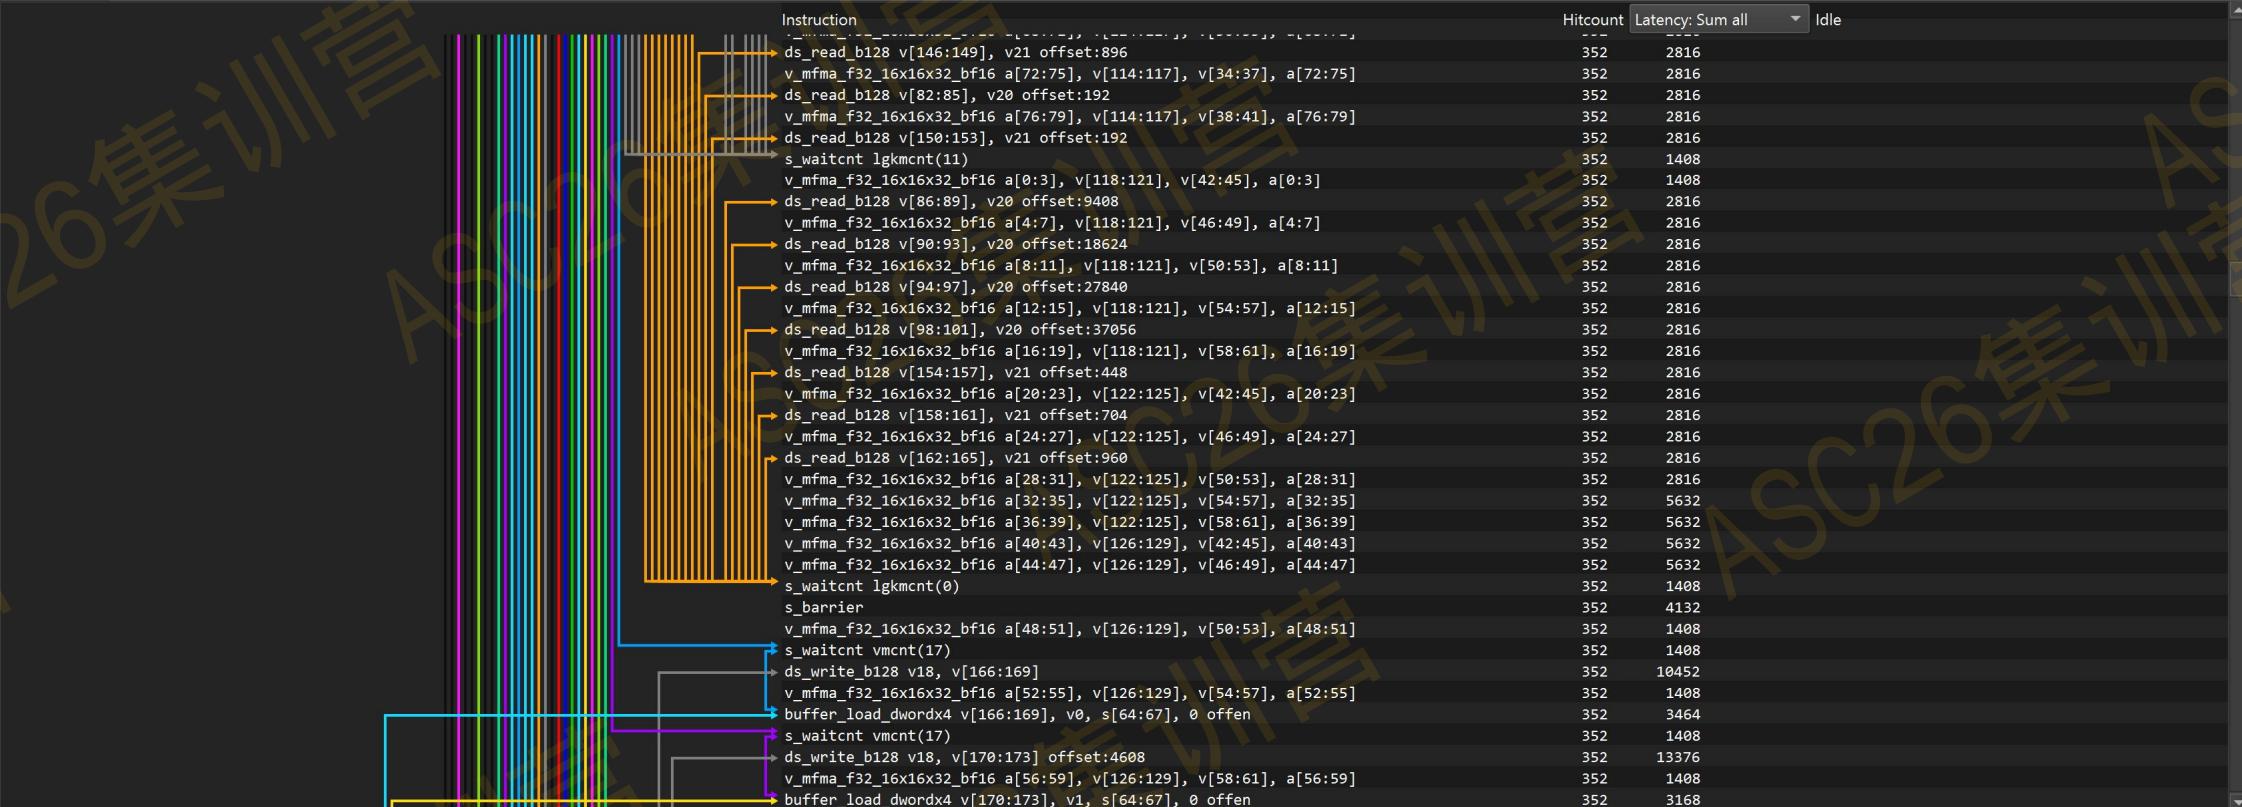
128660

Search Next Prev

History

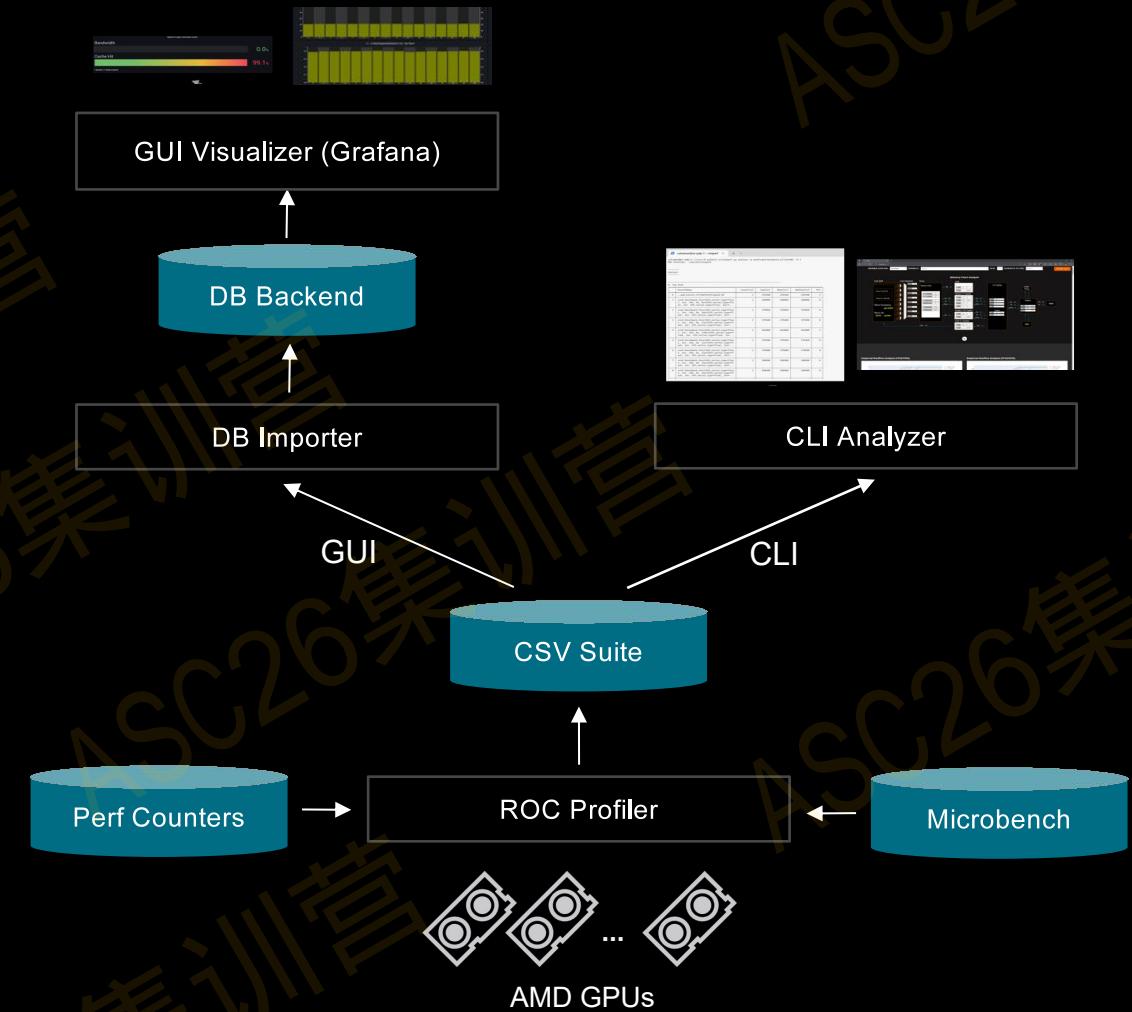
Token Cycle  
1 MATRIX 45976

Instructions Global View Summary Explorer



# The Profiling Tools and Visualization – Omniperf

- Core Omniperf profiler
  - Raw performance counters via application using ROCProfiler
  - Hierarchical roofline data is obtained by a set of micro-benchmarks
- Grafana server for Omniperf
  - Database: Raw performance counters are imported into a MongoDB
  - Grafana GUI: It displays the relevant performance metrics and visualization by retrieving the data from database
- Omniperf Standalone GUI Analyzer
  - Omniperf provides a standalone GUI to enable basic performance analysis without the need to import data into a database instance.
- Features
  - Speed-of-Light (SOL)
  - Hardware Block-level SOL Evaluations
  - Roofline Analysis
  - ...



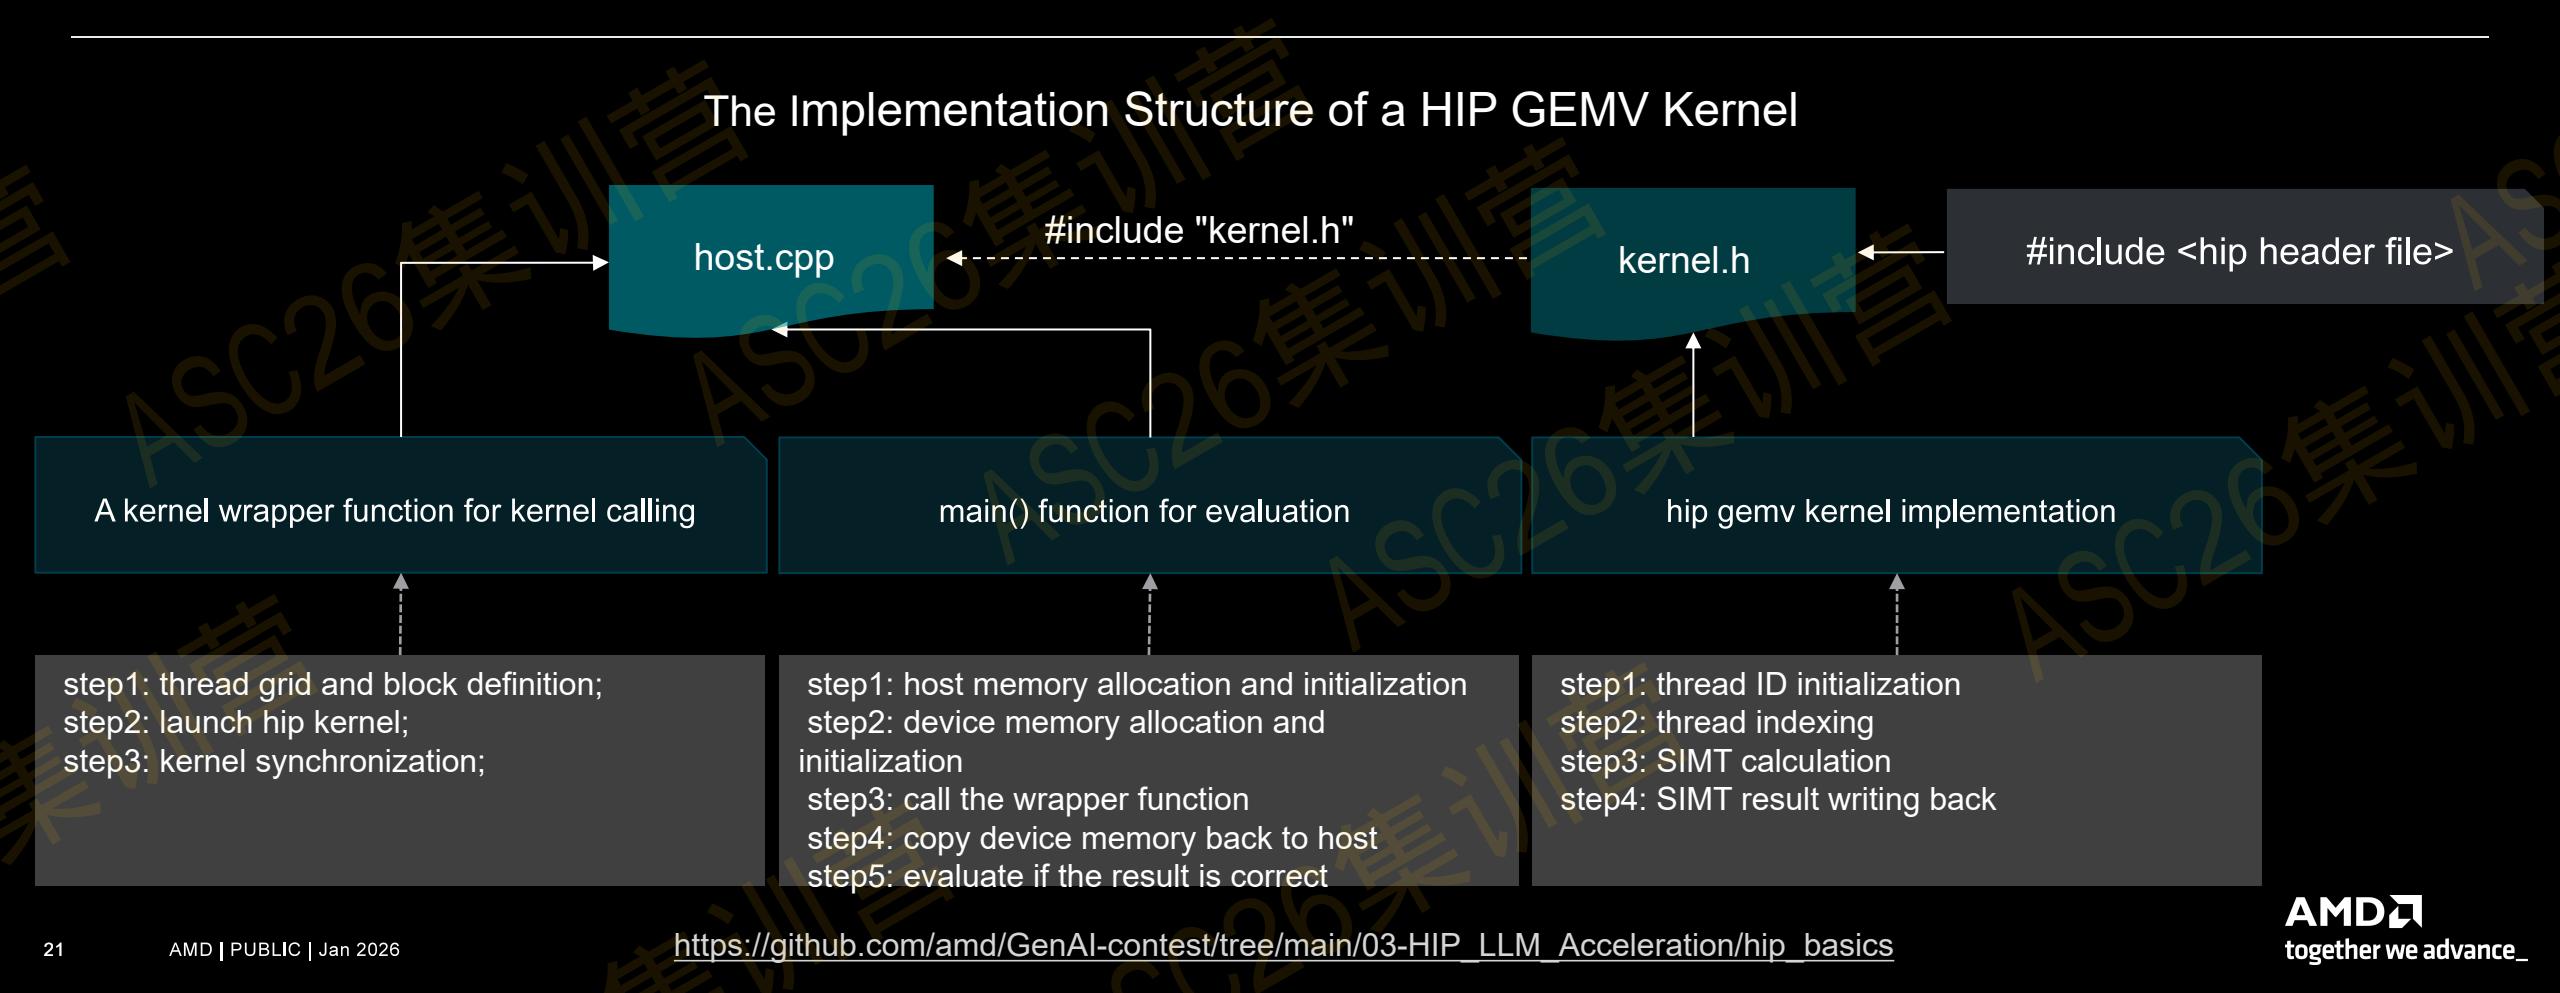
# Agenda

---

1. Introduction to the AMD ROCm™ Software Stack
2. Transitioning Workloads to AMD GPUs
3. Performance Optimization
  - Optimizing application using popular libraries
  - Profiling the models
  - Adding HIP kernel to implement a custom layer
4. Available Collaterals, Q&A

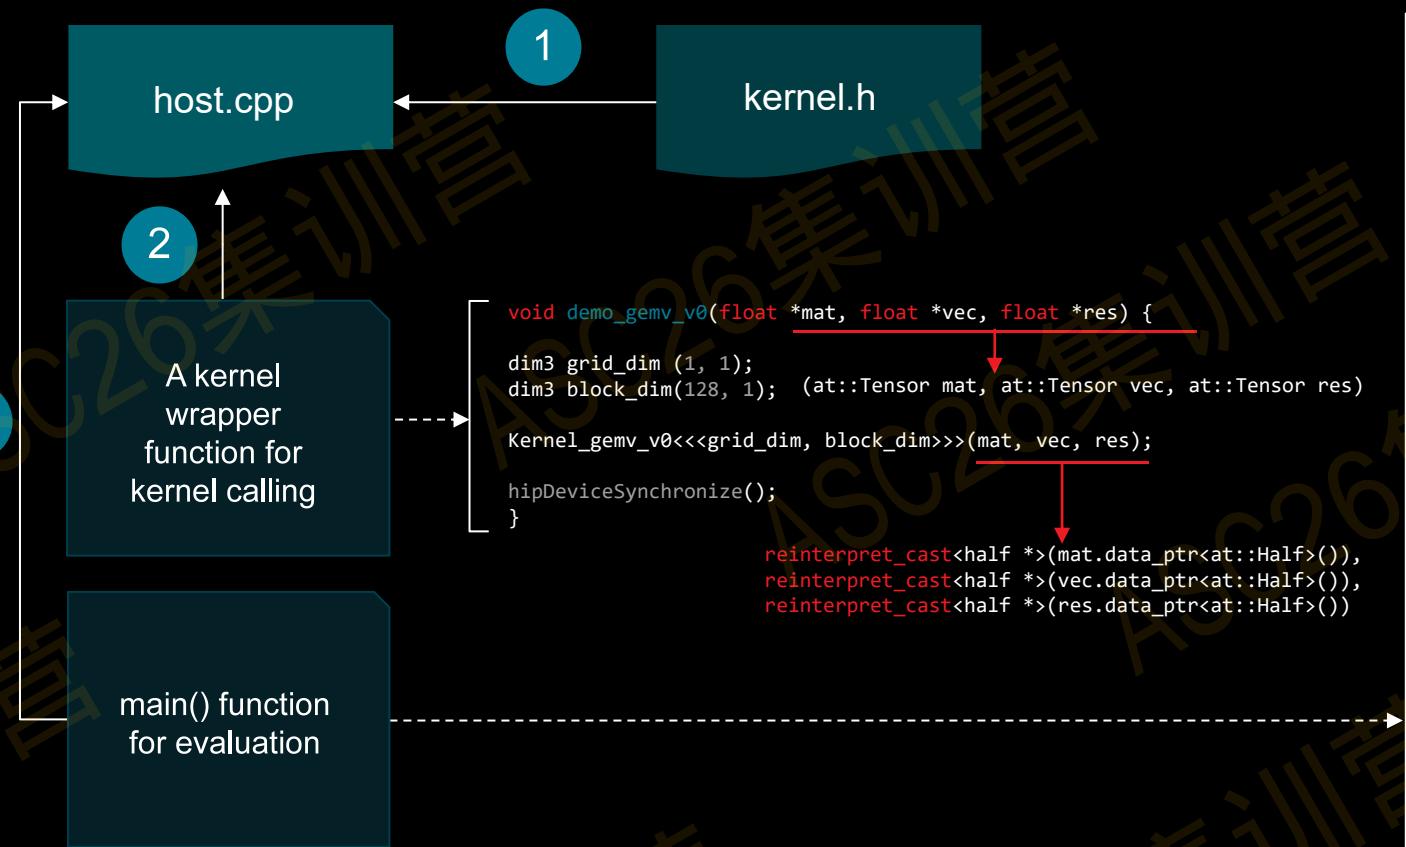
# ROCM Core - Custom HIP GEMV Kernel “hello world” sample

- Given a matrix ( $M \times N$ ), a vector ( $N \times 1$ ), GEMV(matrix, vector) produces an output vector ( $M \times 1$ )
- GPU kernel (kernel.h) launched from host (host.cpp) explores the GPU compute capability by a single instruction multiple threads (SIMT) design



# HIP GEMV Host Code Design

- Given a matrix (128 x 4), a vector (4 x 1), GEMV(matrix, vector) produces an output vector (128 x 1),
- A simple thread parallelism is to employ 128 threads to compute 128 rows in parallel



```

int main() {
    int mat_rows = 128;
    int vec_cols = 4;

    // Allocate memory on CPU
    float* mat = (float*)malloc(sizeof(float) * mat_rows * vec_cols);
    float* vec = (float*)malloc(sizeof(float) * vec_cols);
    float* res = (float*)malloc(sizeof(float) * mat_rows);

    // Fill in some data into mat and vec
    for (int i = 0; i < mat_rows * vec_cols; ++i)
        mat[i] = (float)1.1f;
    for (int i = 0; i < vec_cols; ++i)
        vec[i] = (float)2.2f;

    // Allocate memory on GPU
    float *d_mat, *d_vec, *d_res;
    hipMalloc((void**)&d_mat, mat_rows * vec_cols * sizeof(float));
    hipMalloc((void**)&d_vec, vec_cols * sizeof(float));
    hipMalloc((void**)&d_res, mat_rows * sizeof(float));

    // Host to Device
    hipMemcpy(d_mat, mat, (mat_rows * vec_cols) * sizeof(float),
    hipMemcpyHostToDevice);
    hipMemcpy(d_vec, vec, (vec_cols) * sizeof(float), hipMemcpyHostToDevice);

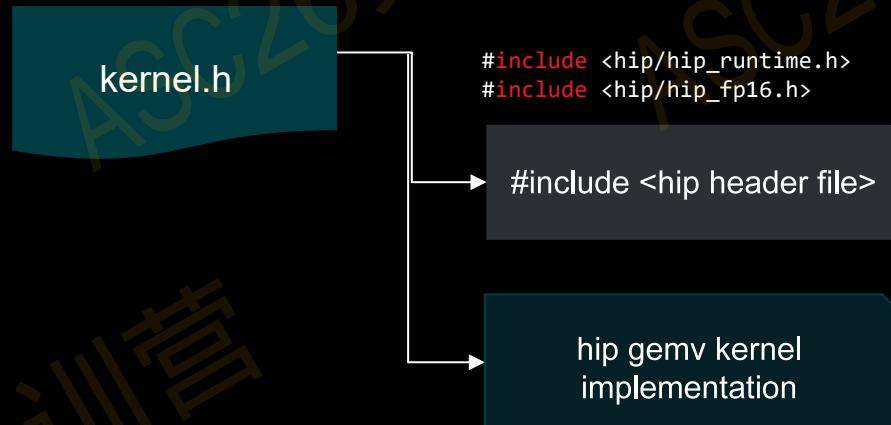
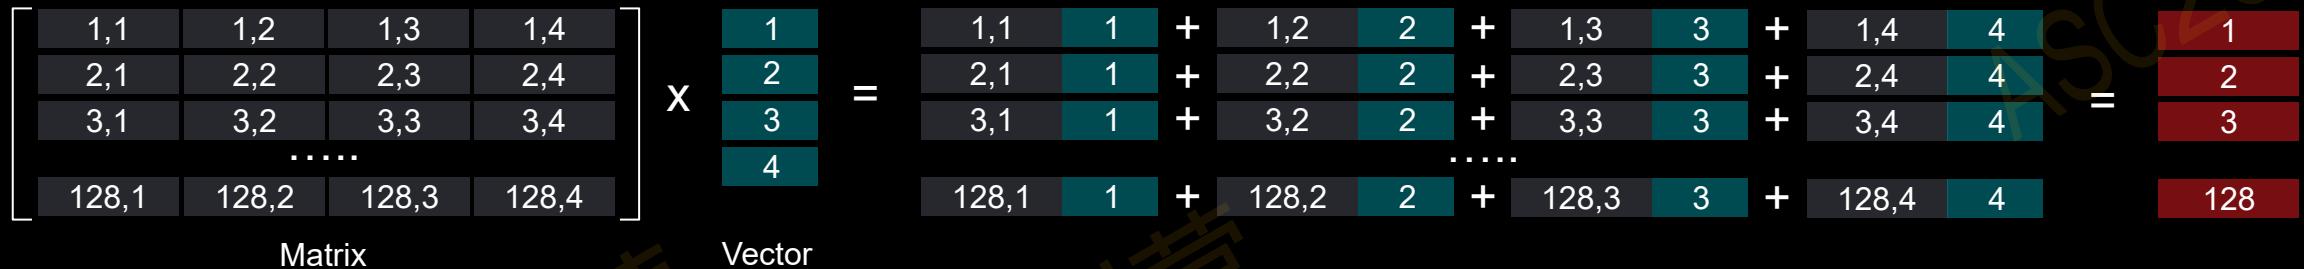
    // Launch kernel
    demo_gemv_v0(d_mat, d_vec, d_res);

    // Device to Host
    hipMemcpy(res, d_res, (mat_rows) * sizeof(float), hipMemcpyDeviceToHost);

    // Print result
    for (int i=0; i< mat_rows; ++i)
        printf("%f ", res[i]);
}

```

# HIP GEMV Kernel Design



```

hipcc --offload-arch=gfx1100 host.cpp -o gemv_v0
./gemv_v0
  
```

```

__global__ void kernel_gemv_v0(float *mat, float *vec, float* res) {
    unsigned int tid = threadIdx.x;
    unsigned int row = tid;
    unsigned int start_idx = 4 * row;

    float mat_h0 = mat[start_idx];
    float mat_h1 = mat[start_idx + 1];
    float mat_h2 = mat[start_idx + 2];
    float mat_h3 = mat[start_idx + 3];

    float vec_h0 = vec[0];
    float vec_h1 = vec[1];
    float vec_h2 = vec[2];
    float vec_h3 = vec[3];

    float sum = 0.0;
    sum += (mat_h0) * (vec_h0);
    sum += (mat_h1) * (vec_h1);
    sum += (mat_h2) * (vec_h2);
    sum += (mat_h3) * (vec_h3);

    res[row] = sum;
}

half *mat, half *vec, half *res)
{
    float mat_h0 = mat[start_idx];
    float mat_h1 = mat[start_idx + 1];
    float mat_h2 = mat[start_idx + 2];
    float mat_h3 = mat[start_idx + 3];

    float vec_h0 = vec[0];
    float vec_h1 = vec[1];
    float vec_h2 = vec[2];
    float vec_h3 = vec[3];

    float sum = 0.0;
    sum += __half2float(mat_h0) * __half2float(vec_h0);
    sum += __half2float(mat_h1) * __half2float(vec_h1);
    sum += __half2float(mat_h2) * __half2float(vec_h2);
    sum += __half2float(mat_h3) * __half2float(vec_h3);

    res[row] = __half2float sum;
}
  
```

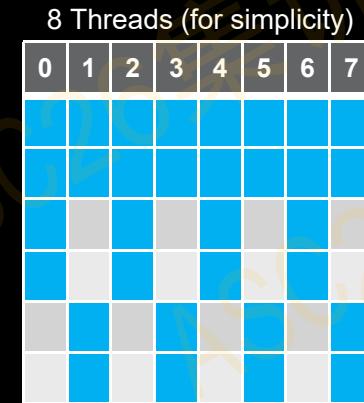
# Performance Optimization – Instruction Throughput

## Control Flow & Divergence

- A wave executes in lockstep. If threads in a wavefront take different branches of an if/else, the GPU executes both paths, masking off threads, leading to divergence and wasted cycles.

Example:

```
int i = blockIdx.x * blockDim.x + threadIdx.x;  
  
if (i % 2 == 0)  
{  
    // half the threads do this  
    out[i] = in[i] * 2.0f;  
}  
else  
{  
    // half the threads do this  
    out[i] = in[i] * 3.0f;  
}
```



## Use Efficient Operations

- Some arithmetic operations are more expensive than others. For example, multiplication is typically faster than division.

## Trade Precision for Speed

- Consider using single-precision arithmetic instead of double-precision if possible.

## Leverage Intrinsic Functions

- Intrinsic functions are predefined functions available in HIP that can often be executed faster than equivalent arithmetic operations.

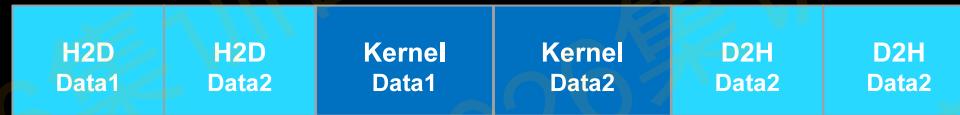
# Performance Optimization – Parallel Execution

## Application Level

- Use asynchronous calls and streams to overlap host/device work. Send serial work to CPU and parallel work to GPU.

Sequential calls:

Default Stream:



Asynchronous calls:

Stream #1:



Stream #2:



## Device Level

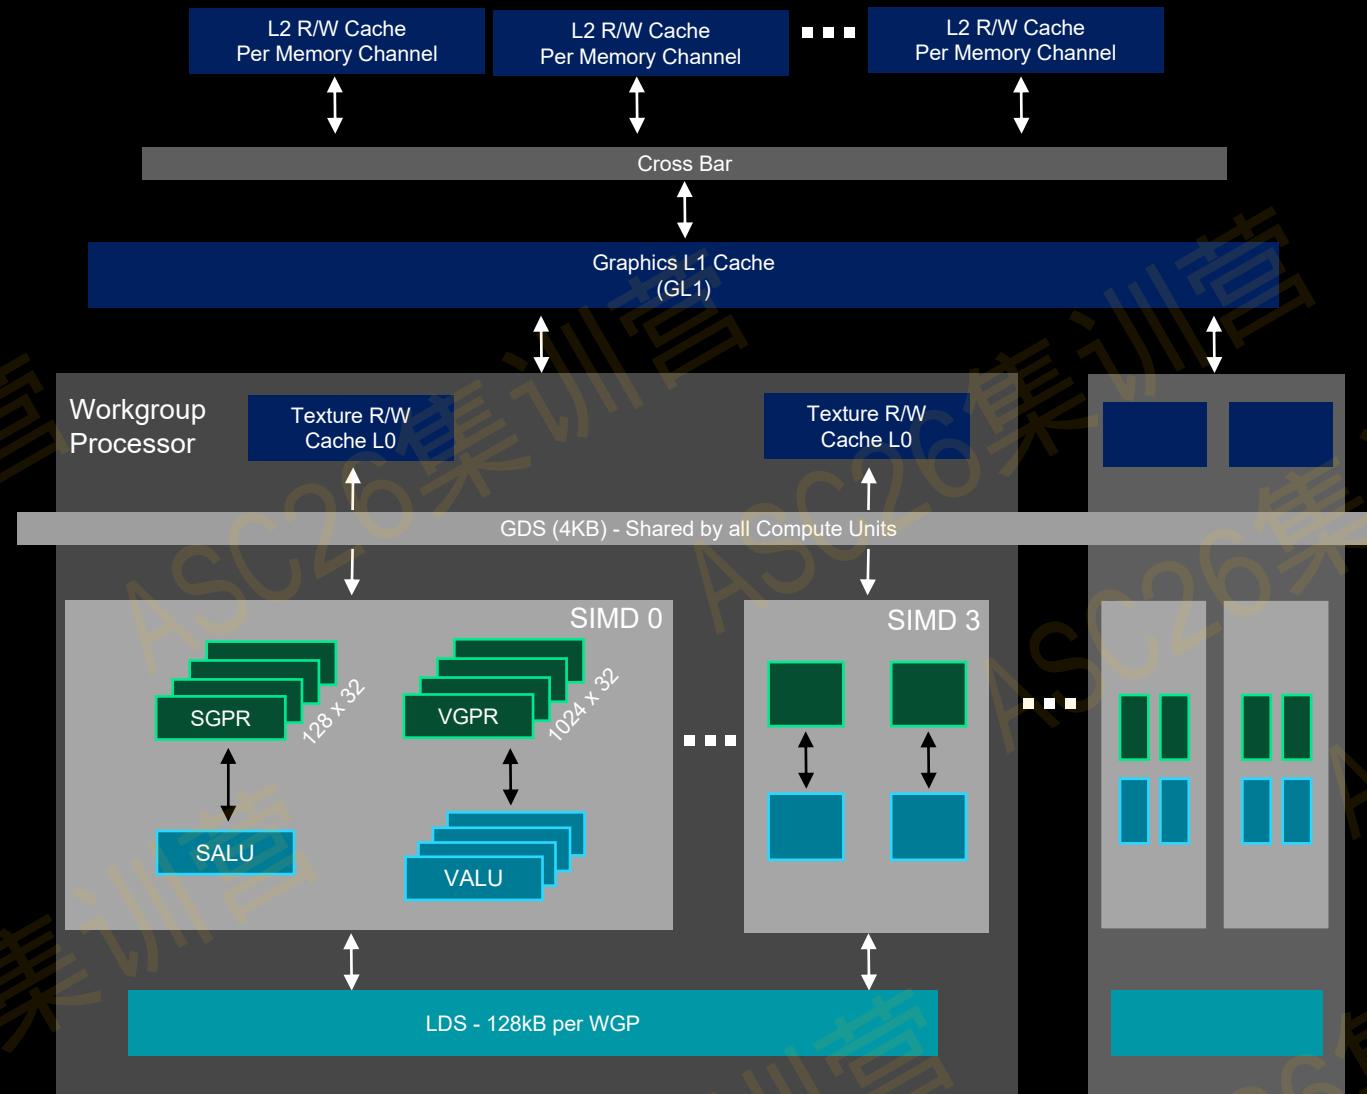
- Maximize utilization by executing enough kernels concurrently while avoiding resource contention.

## Multiprocessor Level

- At its best every clock cycle has an instruction from a warp is ready for execution. This could either be another independent instruction of the same warp or an instruction of another warp.

[https://rocm.docs.amd.com/projects/HIP/en/latest/how-to/performance\\_guidelines.html](https://rocm.docs.amd.com/projects/HIP/en/latest/how-to/performance_guidelines.html)  
[https://rocm.docs.amd.com/projects/HIP/en/latest/how-to/hip\\_runtime\\_api/asynchronous.html](https://rocm.docs.amd.com/projects/HIP/en/latest/how-to/hip_runtime_api/asynchronous.html)

# Performance Optimization – Memory Throughput



## Local Data Share (LDS)

- On-chip shared memory for fast communication and data reuse, often used as a software cache or for cooperative access to off-chip memory.

## Global Data Share (GDS)

- Small on-chip memory shared across all WGP and waves of a kernel. It provides hardware support of append/consume patterns and control data for compute kernels, reduction operations, etc.

## Device Memory Hierarchy (L2 → L1 → L0)

- Multiple L2 cache channels feed read-only L1 and per-WGP L0 caches for off-chip memory accesses. Specialized cache-less load instructions allow direct device memory reads when needed, while caches improve reuse and aggregate scattered accesses.

# Performance Optimization – Memory Throughput

## Local Data Share (LDS)

**Bank Conflict:** It occurs when multiple threads in the same wave access the same bank in shared memory. In this case, accesses get serialized, leading to inferior performance.

$$\text{bank} = \left( \frac{\text{address in bytes}}{4} \right) \bmod 32$$

(Sample: For AMD GCN architecture)

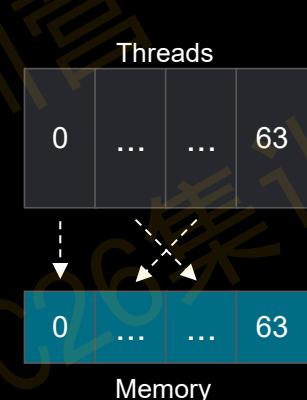
### Optimizations:

- Padding: Change the bank mapping  
`__shared__ float tile[32][33];`
- XOR Preshuffle: Permute the column indices for each row using XOR.
- Use CK Tile abstractions: They automatically handle bank conflict avoidance.
- Consider access patterns: Design algorithms with bank-friendly patterns.

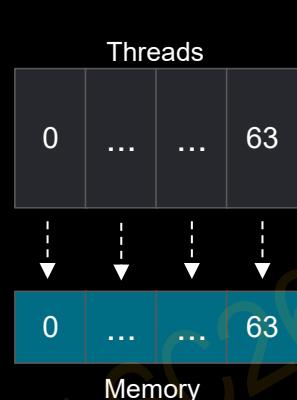
## Device Memory

**Coalescing:** A memory access pattern is coalesced when consecutive threads access consecutive addresses. The hardware can combine them into fewer and wider transactions.

### Uncoalesced Access



### Coalesced Access



### Optimizations:

- Avoid strided access: Array of Structures (AoS) → Structures of Arrays (SoA).
- Align or pad data: Achieve reading/writing contiguous segments.

[https://rocm.docs.amd.com/projects/HIP/en/latest/understand/programming\\_model.html](https://rocm.docs.amd.com/projects/HIP/en/latest/understand/programming_model.html)

[https://rocm.docs.amd.com/projects/composable\\_kernel/en/latest/conceptual/ck\\_tile/hardware/lds\\_bank\\_conflicts.html](https://rocm.docs.amd.com/projects/composable_kernel/en/latest/conceptual/ck_tile/hardware/lds_bank_conflicts.html)

# Agenda

---

1. Introduction to the AMD ROCm™ Software Stack
2. Transitioning Workloads to AMD GPUs
3. Performance Optimization
  - Optimizing application using popular libraries
  - Profiling the models
  - Adding HIP kernel to implement a custom layer
4. Available Collaterals, Q&A

# AMD ROCm™ Software Developer Hub

Initiative to Educate and Increase ROCm™ Software Stack User Base and Adoption

## High-level Overview

Familiarize yourself with the ecosystem  
General introduction of ROCm Software

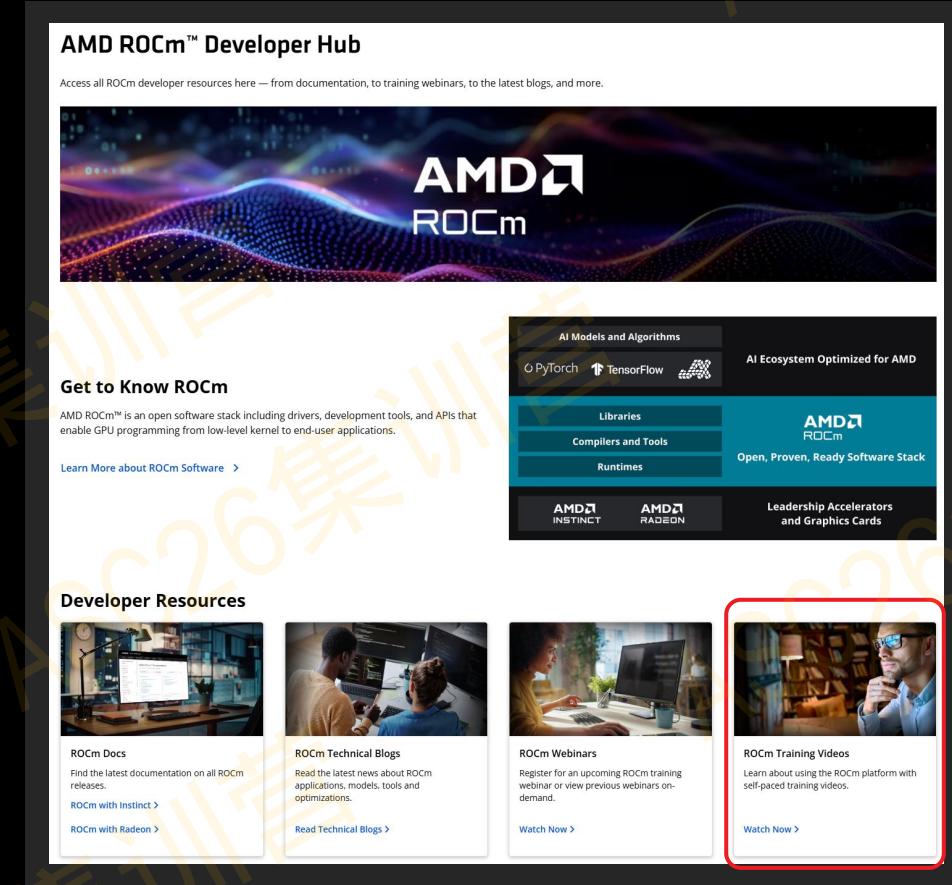
## Increase Understanding

Attend ROCm webinars  
View one of the many training videos

## Build Comprehension

Purchase ROCm textbook  
See the latest news on ROCM blog

ROCM Developer Hub



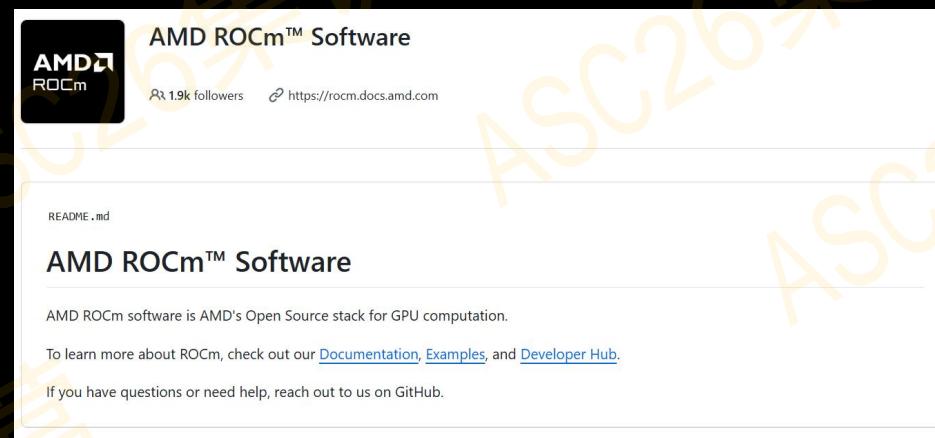
# AMD ROCm™ Documentation & Github Repository

Playground for Professional Developers

## Dive Deeper

Refer to ROCm [documentation](#)

Make contributions to all major components on [Github](#)



## ROCm Github Organization



Join Us

Registration Page



AMD Dev Assistant



**AMD**  
AI Developer Program

<https://account.amd.com/en/forms/registration/ai-dev-program.html>

# Disclaimer and Attributions

## DISCLAIMER

The information contained herein is for informational purposes only, and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale. GD-18

©2025 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, AMD Instinct, AMD ROCm, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.



ASC26集训营