

# WHERE TO POSITION AMD INSTINCT<sup>TM</sup> ACCELERATORS

The AMD Instinct<sup>™</sup> MI200 series accelerators are the newest data center GPUs from AMD, designed to power discoveries in mainstream servers and supercomputers, including some of the largest exascale systems, enabling scientists to tackle our most pressing challenges from climate change to vaccine research.

## AMD INSTINCT ACCELERATORS AND PARTNER PLATFORMS

PARTNER SERVERS	DELL POWEREDGE R7525	DELL POWEREDGE R7515	DELL POWEREDGE R750XA
SERVER FOCUS AREAS (FROM PARTNER)	GENERAL PURPOSE, HPC, ML	GENERAL PURPOSE, BIG DATA, VIRTUALIZATION	AI, ML/DL, HPC, VIRTUALIZATION
AMD INSTINCT GPUS SUPPORTED	UP TO 3X AMD INSTINCT MI210 GPU	UP TO 1X AMD INSTINCT MI210 GPU	UP TO 4X AMD INSTINCT MI210 GPU

AMD INSTINCT <sup>™</sup> MI210 PCIe <sup>®</sup> FASTEST GPU FOR TRADITIONAL PCIE SERVERS	DELIVERED PERFORMANCE	A100 PCIe®	MI210	INSTINCT <sup>™</sup> ADVANTAGE
	FP64	9.7 TF	22.6 TF	2.3X
AMD	FP32	19.5 TF	22.6 TF	1.2X
INSTINCT	FP64 (TENSOR vs. MATRIX)	19.5 TF	45.3 TF	2.3X
	FP32 MATRIX	N/A	45.3 TF	N/A

## **COMMON CONFIGURATIONS AND POSITIONING**

	DEVELOPMENT & POC	INDUSTRY VERTICALS				
Common Configs	ALL VERTICALS	UNIVERSITIES	HPC LABS	AI/ML	GEOSCIENCE	MATERIAL & LIFE SCIENCE
(x2) MI210	AMDA INSTINCT	AMDA INSTINCT				AMDA INSTINCT
(x4) MI210		AMDA INSTINCT			AMDA INSTINCT	
(x8) MI210				AMDA INSTINCT	AMDA INSTINCT	
(x4) MI250			AMDA INSTINCT	AMDA INSTINCT	AMDA INSTINCT	

## AMD INSTINCT<sup>™</sup> MI200 SERIES BENEFITS

#### AMD INSTINCT<sup>®</sup> MI200 SERIES BENEFITS FOR AI & ML

#### Faster time to learning and insights

- MI200 Series accelerators provide up to 383 TFLOPS peak theoretical half-precision (FP16) <sup>MI200-01</sup>
- Up to 1.6x the memory capacity and peak theoretical bandwidth than competitive GPUs MI200-07

#### Train large models at scale

 Leverage highspeed GPU to GPU Infinity Fabric Link @600GB/s to scale your large ML models across multiple GPUs

#### SUPPORTED FRAMEWORKS

**O PyTorch** PyTorch is a GPU accelerated tensor computational framework with a Python front end.

48

### 1 TensorFlow

TensorFlow is an open-source software library for numerical computation using data flow graphs.

#### AMD INSTINCT<sup>®</sup> MI200 SERIES BENEFITS FOR MATERIAL & LIFE SCIENCE

#### Speed time to discovery

 Leadership Peak TFLOPS performance for computebound bioscience workloads MI200-01, MI200-02

## Handle massively complex and data intense simulations

 AMD Instinct<sup>™</sup> MI200 series accelerators offer up to 3.2 TB/s theoretical memory throughput, while NVIDIA Ampere A100 accelerators deliver up to 2<u>TB/s<sup>MI200-07</sup></u>



#### Solve O&G workloads fast

 Leadership single (FP32) and double precision (FP64) peak theoretical performance MI200-02, MI200-03

#### Identify oil reservoir opportunities fast

 AMD Instinct<sup>™</sup> MI250 accelerators offer up to 1.6x the theoretical memory throughput of the highest bandwidth NVIDIA Ampere A100 accelerators <sup>MI200-07</sup>





#### **GEOSCIENCE APPLICATIONS AVAILABLE ON AMD INFINITY HUB**

Globe

SPECFEM3D

Cartesian

MiniMod INSTINCT

and the second second

BRINGING EXASCALE-CLASS TECHNOLOGIES TO MAINSTREAM HPC & AI

AMDA NSTINCT

## CONVERT YOUR CUDA CODE FOR AMD GPUS

EXTEND YOUR APPLICATION PLATFORM SUPPORT BY CONVERTING CUDA $^{\circ}$  CODE TO ROCM

#### **HIPIFY-PERL**

- Easiest to use; point at a directory and it will hipify CUDA code
- Very simple string replacement technique; may require manual postprocessing
- Recommended for quick scans of projects

#### **HIPIFY-CLANG**

- More robust translation of the code
- Generates warnings and assistance for additional analysis
- High quality translation, particularly for cases where the user is familiar with the make system



## ADDITIONAL RESOURCES

PORTING CUDA TO HIP ON DELL POWEREDGE SERVERS

INTRODUCTION TO AMD ROCm<sup>™</sup> ECOSYSTEM: ACCELERATING HPC WORKLOADS WITH AMD INSTINCT<sup>™</sup> GPUS

#### **SPEED TIME TO SCIENCE**

The AMD Infinity Hub contains a collection of advanced GPU software containers and deployment guides for HPC, AI & Machine Learning applications, enabling researchers, scientists and engineers to speed up their time to science **AMD.COM/INFINITYHUB** 



#### START ACCELERATING

AMD Accelerator Cloud offers complimentary access to the latest AMD hardware and software for instant hands-on experience with next-generation capabilities in high performance computing (HPC) and AI. **AMD.COM/AAC** 

# 

he information duration for informational purposes only, and is subject to charge webbaar notes, with sequence of this document, in the sprearation of this doc