AMD RADEON INSTINCT™ MI50
UNLEASH DISCOVERY

The AMD Radeon Instinct™ MI50 server accelerator designed on the world's First 7nm FinFET technology process brings customers a full-feature set based on the industry's newest technologies. The MI50 is AMD's workhorse accelerator offering that is ideal for large scale deep learning. Delivering up to 26.5 TFLOPS of native half-precision (FP16) or up to 13.3 TFLOPS single-precision (FP32) peak floating point performance and INT8 support and combined with 16GB or 32GB of high-bandwidth HBM2 ECC memory, the AMD Radeon Instinct™ MI50 brings customers finely balanced performance needed for enterprise-class, mid-range compute capable of training complex neural networks for a variety of demanding deep learning applications in a cost-effective design

Combine this outstanding performance with the MI50's advanced I/O capabilities, with support for PCIe® Gen 4 technology and the addition of AMD's dual Infinity Fabric™ Link technology providing up to 248 GB/s of aggregate GPU data transfer bandwidth[1]; and the advanced efficiencies achieved with the use of AMD's 2nd Gen "Vega" 7nm technology; and you truly have one of the markets' most advanced, working-class server GPUs available for a variety of deep learning, cloud and general purpose HPC.

Optimized Deep Learning Operations

Optimized Deep Learning Operations with flexible mixed-precision **FP16, FP32, INT8 & INT4** capabilities bring customers supercharged compute performance to meet today's demanding system requirements of handling large data efficiently for training complex neural networks and running inference against those networks used in deep learning.

AMD Infinity Fabric Link

Two Infinity Fabric Links per GPU for high speed Direct-Connect GPU hives delivering up to **184 GB/s** GPU theoretical peer-to-peer bandwidth – 4.75x faster than using PCIe 3.0 alone.[1]

Rapid FP64 Performance for HPC Workloads

The Radeon Instinct MI50 delivers great double precision performance with up to **6.6** TFLOPS FP64 performance, enabling scientists and researchers across the globe to more efficiently process HPC parallel codes across several industries including life sciences, energy, finance, automotive and aerospace, academics, government, defense and more.

Ultra-Fast HBM2 Memory

With 16GB or 32GB HBM2 memory utilizing a four-stack memory configuration, the new **AMD Radeon Instinct MI50** delivers ultra-high memory bandwidth with up to **1 TB/s** memory bandwidth.

AMD Radeon Instinct MI50 Ready Servers

PENGUIN COMPUTING	GIGABYTE™			
Product: Altus XE4218GT **Form Factor:** 4U/2P/8GPU/PCIe Gen 4 + Switch **Specs:** 2x 2nd Gen AMD EPYC™ Processors + 8x AMD Radeon Instinct™ MI50 GPU accelerators	**Product:** G292-Z20 **Form Factor:** 2U/1P/8GPU/PCIe Gen 4 + Switch **Specs:** 1x 2nd Gen AMD EPYC™ Processor + 8x AMD Radeon Instinct™ MI50 GPU accelerators	**Product:** G292-Z40 **Form Factor:** 2U/2P/8GPU/PCIe Gen 4 + Switch **Specs:** 2x 2nd Gen AMD EPYC™ Processors + 8x AMD Radeon Instinct™ MI50 GPU accelerators	**Product:** G482-Z52 **Form Factor:** 4U/2P/8GPU/PCIe Gen 4 **Specs:** 2x 2nd Gen AMD EPYC™ Processors + 8x AMD Radeon Instinct™ MI50 GPU accelerators	**Product:** G482-Z51 **Form Factor:** 4U/2P/8GPU/PCIe Gen 4 **Specs:** 2x 2nd Gen AMD EPYC™ Processors + 8x AMD Radeon Instinct™ MI50 GPU accelerators

AMD

ROCm Open Ecosystem

The ROCm Open eCosystem, is an open-source HPC/Hyperscale-class platform for GPU computing that's also programming-language independent. ROCm brings choice, minimalism and modular software development to GPU computing.

Deep Learning and HPC Applications

Cloud / Hyperscale

Financial Services

Energy

Life Sciences

Automotive

HPC

PYTORCH   TensorFlow

Open Programing Languages
OpenMP, HIP, OpenCL and Python

Key Features

PERFORMANCE

Compute Units	60
Stream Processors	3,840
Peak INT8	Up to 53.6 TOPS
Peak FP16	Up to 26.5 TFLOPS
Peak FP32	Up to 13.3 TFLOPS
Peak FP64	Up to 6.6 TFLOPS
Bus Interface	PCIe®Gen 3 and Gen 4 Supported[2]

MEMORY

Memory Size	16GB or 32GB HBM2
Memory Interface	4,096-Bits
Memory Clock	1 GHz
Memory Bandwidth	Up to 1 TB/s

RELIABILITY

ECC (Full-chip)	Yes[3]
RAS Support	Yes[4]

SCALABILITY

Infinity Fabric Links	2
OS Support	Linux® 64-bit
ROCm Compatible	Yes

BOARD DESIGN

Board Form Factor	Full-Height, Dual Slot
Length	10.5" Long
Thermal	Passively Cooled
Max Power	300W TDP
Warranty	Three Year Limited[5]

For More Information Visit:
AMD.com/INSTINCT, AMD.com/ROCm and AMD.com/HPC

AMD