

THE EXASCALE ERA IS HERE


AMD
INSTINCT

AMD INSTINCT™ MI200 SERIES ACCELERATOR

World's Fastest HPC and AI Accelerator¹

The era of exascale is here. Immense computational power coupled with the fusion of HPC and AI is enabling researchers and scientists to tackle our most pressing challenges from climate change to vaccine research. With the AMD Instinct™ MI200 accelerators and ROCm™ 5.0 software ecosystem, innovators can tap the power of the world's most powerful HPC and AI data center GPUs to accelerate their time to science and discovery.¹

Based on the 2nd Gen AMD CDNA™ architecture, AMD Instinct™ MI200 accelerators deliver a quantum leap in HPC and AI performance over competitive data center GPUs today. With an up to 4x advantage in HPC performance compared to competitive GPUs, the MI200 accelerator is the first data center GPU to deliver 383 teraflops of theoretical mixed precision FP16 performance for deep learning training, offering users a powerful platform to fuel the convergence of HPC and AI.¹



Graph 1: Peak TFLOPS across range of mixed-precision Compute¹

Innovations Delivering Performance Leadership

AMD innovations in architecture, packaging and integration are pushing the boundaries of computing by unifying the most important processors in the data center, the CPU and the GPU accelerator. With industry-first multi-chip GPU modules along with 3rd Gen AMD Infinity Architecture, AMD is delivering performance, efficiency and overall system throughput for HPC and AI using AMD EPYC™ CPUs and AMD Instinct™ MI200 series accelerators



Key Features

PERFORMANCE	MI250	MI250X
Compute Units	208CU	220CU
Stream Processors	13,312	14,080
Peak FP64/FP32 Vector	45.3 TFLOPS	47.9 TFLOPS
Peak FP64/FP32 Matrix	90.5 TFLOPS	95.7 TFLOPS
Peak FP16/BF16	362.1 TFLOPS	383.0 TFLOPS
Peak INT4/INT8	362.1 TOPS	383.0 TOPS

MEMORY

Memory Size	128GB HBM2e	128GB HBM2e
Memory Interface	8,192 bits	8,192 bits
Memory Clock	1.6GHz	1.6GHz
Memory Bandwidth	up to 3.2TB/sec ²	up to 3.2TB/sec ²

RELIABILITY	MI250	MI250X
ECC (Full-chip)	Yes	Yes
RAS Support	Yes	Yes

SCALABILITY

Infinity Fabric™ Links	up to 6	up to 8
Coherency Enabled	No	Yes
OS Support	Linux™ 64 Bit	Linux 64 Bit
AMD ROCm™ Compatible	Yes	Yes

BOARD DESIGN

Form Factor	OAM	OAM
Thermal	Passive & Liquid	Passive & Liquid
Max Power	500W & 560W TDP	500W & 560W TDP
Bus Interface	PCIe® Gen 4 Support	
Warranty	3 Year Limited ⁴	

Алматы (7273)495-231
Ангарск (3955)60-70-56
Архангельск (8182)63-90-72
Астрахань (8512)99-46-04
Барнаул (3852)73-04-60
Белгород (4722)40-23-64
Благовещенск (4162)22-76-07
Брянск (4832)59-03-52
Владивосток (423)249-28-31
Владикавказ (8672)28-90-48
Владимир (4922)49-43-18
Волгоград (844)278-03-48
Вологда (8172)26-41-59
Воронеж (473)204-51-73
Екатеринбург (343)384-55-89

Иваново (4932)77-34-06
Ижевск (3412)26-03-58
Иркутск (395)279-98-46
Казань (843)206-01-48
Калининград (4012)72-03-81
Калуга (4842)92-23-67
Кемерово (3842)65-04-62
Киров (8332)68-02-04
Коломна (4966)23-41-49
Кострома (4942)77-07-48
Краснодар (861)203-40-90
Красноярск (391)204-63-61
Курск (4712)77-13-04
Курган (3522)50-90-47
Липецк (4742)52-20-81
Россия +7(495)268-04-70

Магнитогорск (3519)55-03-13
Москва (495)268-04-70
Мурманск (8152)59-64-93
Набережные Челны (8552)20-53-41
Нижний Новгород (831)429-08-12
Новокузнецк (3843)20-46-81
Ноябрьск (3496)41-32-12
Новосибирск (383)227-86-73
Омск (3812)21-46-40
Орел (4862)44-53-42
Оренбург (3532)37-68-04
Пенза (8412)22-31-16
Петрозаводск (8142)55-98-37
Псков (8112)59-10-37
Пермь (342)205-81-47
Казахстан +7(7172)727-132

Ростов-на-Дону (863)308-18-15
Рязань (4912)46-61-64
Самара (846)206-03-16
Санкт-Петербург (812)309-46-40
Саратов (845)249-38-78
Севастополь (8692)22-31-93
Саранск (8342)22-96-24
Симферополь (3652)67-13-56
Смоленск (4812)29-41-54
Сочи (862)225-72-31
Ставрополь (8652)20-65-13
Сургут (3462)77-98-35
Сыктывкар (8212)25-95-17
Тамбов (4752)50-40-97
Тверь (4822)63-31-35
Киргизия +996(312)96-26-47

Тольятти (8482)63-91-07
Томск (3822)98-41-53
Тула (4872)33-79-87
Тюмень (3452)66-21-18
Ульяновск (8422)24-23-59
Улан-Удэ (3012)59-97-51
Уфа (347)229-48-12
Хабаровск (4212)92-98-04
Чебоксары (8352)28-53-07
Челябинск (351)202-03-61
Череповец (8202)49-02-64
Чита (3022)38-34-83
Якутск (4112)23-90-97
Ярославль (4852)69-52-93

<https://amd.nt-rt.ru/> || aid@nt-rt.ru

Ecosystem without Borders

AMD ROCm™ is an open software platform allowing researchers to tap the power of AMD Instinct™ accelerators to drive scientific discoveries. The ROCm platform is built on the foundation of open portability, supporting environments across multiple accelerator vendors and architectures. With ROCm 5.0, AMD extends its platform powering top HPC and AI applications with AMD Instinct MI200 series accelerators, increasing accessibility of ROCm for developers and delivering outstanding performance across key workloads.

HPC and MACHINE LEARNING APPLICATIONS



HPC



Life Sciences



Chemistry



Energy



Weather



Astrophysics



Automotive

Reinforcement
LearningImage | Object |
Video Detection &
Classification

OPEN PROGRAMING WITH CHOICE

OpenMP | HIP | OpenCL™ | Python

OPEN FRAMEWORKS

PyTorch | TensorFlow | ONNX | Kokkos | RAJA

OPTIMIZED LIBRARIES

BLAS | FFT | RNG | SPARSE | THRUST | MIOpen | RCCL

PROGRAMMER AND SYSTEM TOOLS

Debuggers | Performance Analysis | System Management

2nd Generation AMD CDNA™ Architecture

The AMD Instinct™ MI200 series accelerator brings customers the compute engine selected for the first U.S. Exascale supercomputer. Powered by the 2nd Generation AMD CDNA™ architecture, the MI200 series accelerators deliver a quantum leap in HPC and AI performance over competitive data center GPUs today. The AMD Instinct MI200 series GPU delivers industry-leading double precision performance for HPC workloads with up to 47.9TFLOPS peak FP64 performance, enabling scientists and researchers across the globe to process HPC parallel codes more efficiently across several industries.¹

AMD's Matrix Core technology delivers a full range of mixed precision operations bringing you the ability to work with large models and enhance memory-bound operation performance for whatever combination of AI and machine learning workloads you need to deploy. The MI200 offers optimized BF16, INT4, INT8, FP16, FP32, and FP64 Matrix capabilities bringing you supercharged compute performance to meet all your AI system requirements. The AMD Instinct MI200 accelerator handles large data efficiently for training and is the first data center GPU to deliver 383 teraflops of peak FP16 performance for deep learning training.¹

AMD Infinity Fabric™ Link Technology

AMD Instinct MI200 series OAM accelerators with advanced peer-to-peer I/O connectivity through a maximum of eight AMD Infinity Fabric™ links deliver up to 800 GB/s I/O bandwidth performance.³ With a cache coherency enabled solution using Optimized 3rd Gen AMD EPYC™ CPU and MI250X accelerators, Infinity Fabric unlocks the promise of unified computing, enabling a quick and simple on-ramp for CPU codes to accelerated platforms.

Ultra-Fast HBM2e Memory

The AMD Instinct™ MI200 accelerators provide up to 128GB High-bandwidth HBM2e memory with ECC support at a clock rate of 1.6 GHz. and deliver an ultra-high 3.2 TB/s of memory bandwidth to help support your largest data sets and eliminate bottlenecks in moving data in and out of memory.² Combine this performance with the MI200's advanced I/O capabilities and you can push workloads closer to their full potential.

Алматы (7273)495-231
Ангарск (3955)60-70-56
Архангельск (8182)63-90-72
Астрахань (8512)99-46-04
Барнаул (3852)73-04-60
Белгород (4722)40-23-64
Благовещенск (4162)22-76-07
Брянск (4832)59-03-52
Владивосток (423)249-28-31
Владикавказ (8672)28-90-48
Владимир (4922)49-43-18
Волгоград (844)278-03-48
Вологда (8172)26-41-59
Воронеж (473)204-51-73
Екатеринбург (343)384-55-89

Иваново (4932)77-34-06
Ижевск (3412)26-03-58
Иркутск (395)279-98-46
Казань (843)206-01-48
Калининград (4012)72-03-81
Калуга (4842)92-23-67
Кемерово (3842)65-04-62
Киров (8332)68-02-04
Коломна (4966)23-41-49
Кострома (4942)77-07-48
Краснодар (861)203-40-90
Красноярск (391)204-63-61
Курск (4712)77-13-04
Курган (3522)50-90-47
Липецк (4742)52-20-81
Россия +7(495)268-04-70

Магнитогорск (3519)55-03-13
Москва (495)268-04-70
Мурманск (8152)59-64-93
Набережные Челны (8552)20-53-41
Нижний Новгород (831)429-08-12
Новокузнецк (3843)20-46-81
Ноябрьск (3496)41-32-12
Новосибирск (383)227-86-73
Омск (3812)21-46-40
Орел (4862)44-53-42
Оренбург (3532)37-68-04
Пенза (8412)22-31-16
Петрозаводск (8142)55-98-57
Псков (8112)59-10-37
Пермь (342)205-81-47
Казахстан +7(7172)727-132

Ростов-на-Дону (863)308-18-15
Рязань (4912)46-61-64
Самара (846)206-03-16
Санкт-Петербург (812)309-46-40
Саратов (845)249-38-78
Севастополь (8692)22-31-93
Саранск (8342)22-96-24
Симферополь (3652)67-13-56
Смоленск (4812)29-41-54
Сочи (862)225-72-31
Ставрополь (8652)20-65-13
Сургут (3462)77-98-35
Сыктывкар (8212)25-95-17
Тамбов (4752)50-40-97
Тверь (4822)63-31-35
Киргизия +996(312)96-26-47

Тольятти (8482)63-91-07
Томск (3822)98-41-53
Тула (4872)33-79-87
Тюмень (3452)66-21-18
Ульяновск (8422)24-23-59
Улан-Удэ (3012)59-97-51
Уфа (347)229-48-12
Хабаровск (4212)92-98-04
Чебоксары (8352)28-53-07
Челябинск (351)202-03-61
Череповец (8202)49-02-64
Чита (3022)38-34-83
Якутск (4112)23-90-97
Ярославль (4852)69-52-93

<https://amd.nt-rt.ru/> || aid@nt-rt.ru