



Cognitar

PER ASPERA AD ASTRA

Новое поколение вычислительного оборудования для искусственного интеллекта в научных исследованиях

Валерий Егоршев, Когнитар

Июль 2025, Москва



ЦОД ОИЯИ (Дубна)

A photograph of a data center aisle. On the left, a large black server cabinet with double glass doors and green handles is prominent. The APC logo is visible on the top left of the door. To its right, a long row of similar black server racks extends into the distance. The floor is light gray tile. The ceiling has a grid of fluorescent lights and two large white air conditioning units. A yellow and black striped safety line runs along the base of the left wall. A doorway is visible at the end of the aisle on the right.

ЦОД ПИК ПИЯФ ОИЯИ (Гатчина)

Новое поколение процессоров

Е-ядра и Р-ядра в x86 – новая реальность ввиду пришествия ARM



AMD Bergamo 9754 (128 – Е-ядер x86)




Ampere Altra Max (128 ядер ARM)



Процессоры AMD

AMD EPYC™ Processors

Five generations of on time technology innovation



All roadmaps are subject to change.

Процессоры AMD - система обозначений

AMD EPYC™ 9005 Series

Processor Naming Convention

EPYC™ 9535P CPU

Product Family

★ Product Series 9005

Compute

- "F" = High Frequency
- "P" = 1P Capable Only

Generation

Core Count

- Indicates Core Count within the series


100s Digit	0	1	2	3	4	5	6	7	8	9
Cores	8	16	24	32 - 36	48	64 - 72	96	128	144 - 160	192

Performance

- 10s digit – Perf within Core Count
- 8, 9 = reserved
- 7, 6, 5, 4, 3, 2, 1
- Relative Performance within core count
- Higher number = higher perf

Процессоры AMD Turin (Zen 5)

Процессоры AMD Turin (Zen5)



Previewing Today at Computex 2024: **"Turin"**

5th Gen AMD EPYC™

World's best data center CPU.



Architecture

Up to **192** Cores
Up to **384** Threads

SP5 Socket
Compatible with "Genoa"

Available
2H 2024

Процессоры AMD Turin (Zen5)



Процессоры AMD Turin (Zen5)

AMD EPYC™ 9005 Series Processors



Increased core density



Energy efficient



Broad OPN stack

Cores	AMD EPYC	CCD (Zen5/Zen5c)	Base/Boost* (up to GHz)	Default TDP (W)	L3 Cache (MB)	Price (1 KU, USD)
192 cores	9965	"Zen5c"	2.25 / 3.7	500W	384	\$14,813
160 cores	9845	"Zen5c"	2.1 / 3.7	390W	320	\$13,564
144 cores	9825	"Zen5c"	2.2 / 3.7	390W	384	\$13,006
128 cores	9755	"Zen5"	2.7 / 4.1	500W	512	\$12,984
	9745	"Zen5c"	2.4 / 3.7	400W	256	\$12,141
96 cores	9655	"Zen5"	2.6 / 4.5	400W	384	\$11,852
	9655P	"Zen5"	2.6 / 4.5	400W	384	\$10,811
	9645	"Zen5c"	2.3 / 3.7	320W	384	\$11,048
72 cores	9565	"Zen5"	3.15 / 4.3	400W	384	\$10,486
64 cores	9575F	"Zen5"	3.3 / 5.0	400W	256	\$11,791
	9555	"Zen5"	3.2 / 4.4	360W	256	\$9,826
	9555P	"Zen5"	3.2 / 4.4	360W	256	\$7,983
	9535	"Zen5"	2.4 / 4.3	300W	256	\$8,992
48 cores	9475F	"Zen5"	3.65 / 4.8	400W	256	\$7,592
	9455	"Zen5"	3.15 / 4.4	300W	192	\$5,412
	9455P	"Zen5"	3.15 / 4.4	300W	192	\$4,819
36 cores	9365	"Zen5"	3.4 / 4.3	300W	256	\$4,341
32 cores	9375F	"Zen5"	3.8 / 4.8	320W	256	\$5,306
	9355	"Zen5"	3.55 / 4.4	280W	256	\$3,694
	9355P	"Zen5"	3.55 / 4.4	280W	256	\$2,998
	9335	"Zen5"	3.0 / 4.4	210W	256	\$3,178
24 cores	9275F	"Zen5"	4.1 / 4.8	320W	256	\$3,439
	9255	"Zen5"	3.25 / 4.3	200W	128	\$2,495
16 cores	9175F	"Zen5"	4.2 / 5.0	320W	512	\$4,256
	9135	"Zen5"	3.65 / 4.3	200W	64	\$1,214
	9115	"Zen5"	2.6 / 4.1	125W	64	\$726
8 cores	9015	"Zen5"	3.6 / 4.1	125W	64	\$527


Процессоры AMD Turin (Zen5)

“Turin” Continues to Deliver Technology Leadership

Scale-Up

Up to


16 “Zen 5” CCDs
128 Cores / 256 Threads



Scale-Out

Up to

12 “Zen 5c” CCDs
192 Cores / 384 Threads



Consistent features,
ISA, & IPC uplift

SP5 Socket
“Genoa” Compatible

8 to 192 Cores
155W to 500W

Up to
12Ch DDR5-6400*
128 PCIe 5.0/CXL 2.0

Confidential Compute
with Trusted I/O

Процессоры AMD Turin (Zen5)


5th Gen AMD EPYC™ Generational Innovations

Compute

- “Zen5” up to **128 cores** / 256 threads
- “Zen5c” up to **192 cores** / 384 threads
- AVX-512 with **full 512b data path**
- New **500W** performance option
- Faster **5GHz** options
- **3/4nm** Zen cores

I/O & Platform

- 2P and 1P Configurations
- Up to 160 lanes of PCIe® Gen5
- **PCIe link encryption**
- SP5 Compatible with “**Genoa**”
- **CXL® 2.0¹**



Memory

- 12 ch. DDR5 ECC **up to 6400* MT/s**
- Up to 2 DIMMs/channel capacity delivering up to **6TB/socket**
- **Dynamic Post Package Repair (PPR)** for x4 and x8 ECC RDIMMs

Security

- Hardware Root-of-Trust
- **Trusted I/O**
- **FIPS 140-3 in process**

1- CXL Type 162 devices and PCIe link encryption support dependent upon ecosystem readiness
*Standard roadmap offerings on AMD.com support 6000 MT/s
See endnotes 9xx5-048, 072, 083, CD-183A

Процессоры AMD Turin (Zen5)

5th Gen AMD EPYC™ SoC


Blue text indicates significant update from "Zen 4"

Compute

- AMD "Zen5/Zen5c" x86 cores
 - "Zen5" up to 16 CCDs / 128 cores / 256 threads
 - "Zen5c" up to 12 CCDs / 192 cores / 384 threads
- 1MB L2/Core, Up to 32MB L3 per CCD
- ISA Updates: AVX-512 with 512b data path
- Infinity Fabric: 3rd-gen MCM (32+Gbps die-to-die)
- Dynamic PPR for x4 and x8 ECC RDIMMs
- BMC MCA Crash-Dump
- MCA over APML

SP5 Platform

- BIOS Update Required
- 2P and 1P Configurations
- Up to 4 links of 32Gbps Gen 3 Infinity Fabric™
- Compatible with "Genoa IRM" Groups up to 400W
- 500W IRM option for maximum performance
- Flexible topology options
- Server Controller Hub (USB, UART, SPI, I2C, etc.)



Memory

- 12 channel DDR5 with ECC up to 6000 MT/s
- Support for 2, 4, 6, 8, 10, 12 channel memory interleaving
- RDIMM, 3DS RDIMM
- Up to 2 DIMMs/channel capacity of 6TB/socket (256GB 3DS RDIMMs)

I/O

- 2P: up to 160 lanes of PCIe Gen5 with speeds up to 32Gbps/lane
 - up to 12 bonus PCIe Gen3 lanes
- 1P: 128 lanes of PCIe Gen5 with speeds up to 32Gbps/lane
 - up to 8 bonus PCIe Gen3 lanes
- PCIe bifurcations supported: x16, x8, x4, x2 and x1
- PCIe link encryption²
- CXL® 2.0, 4 x16 Capable "P" links; Type 3, Type 1², Type 2² PoC
- Up to 32 IO lanes for SATA
- SDCl (Smart Data Cache Injection)

Security¹

- Trusted IO
- Dedicated Security Subsystem with enhancements
- Hardware Root-of-Trust; Ciphertext-hiding capability


1 - AMD Infinity Guard features vary by EPYC™ Processor generations and/or series. Infinity Guard security features must be enabled by server OEMs and/or Cloud Service Providers to operate. Check with your OEM or provider to confirm support of these features. Learn more about Infinity Guard at <https://www.amd.com/en/technologies/infinity-guard>. GD-183A.

2 - CXL Type 1&2 devices and PCIe link encryption support dependent upon ecosystem readiness.


Процессоры AMD Turin (Zen5)

5th Gen AMD EPYC™ SoC Platform Overview

2P Configuration



1P Configuration



**Strong upgrade to existing SP5
(4th Gen AMD EPYC™ CPU)
in the same platforms**

- 25% improvement in DRAM speed (4800 → up to 6000) using JEDEC-standard (non-proprietary) DIMMs
- 1P PCIe® aggregate bandwidth improvement due to internal SOC topology changes
- Enhanced platform option for 500W TDP capability

Процессоры AMD Turin (Zen5)

CXL® Overview

CXL industry standard Cache-Coherent Interconnect for Processors, Memory Expansion and Accelerators

5th Gen AMD EPYC™ CPU supports CXL® 2.0

- Types 1, 2 & 3 *


Additional CXL Capabilities

- Tiered memory support
 - Multiple CXL devices combined into single interleaved NUMA node
 - Headless NUMA nodes
 - "Pinned"-memory and Secure Memory datamover
 - Memory profiling by HV/Guest (IBS filtering, IBS virtualization)
- x86 QoS support
 - Separate x86 MBM BW limits for DRAM and CXL memory
- AMD SEV/SNP features fully supported on CXL Type 3 memory

Low latency memory attach

- Device-dependent; similar to remote 2P socket latency for local CXL-attach memory

5th Gen AMD EPYC™ CPU Focus



- CXL.io: Discovery, configuration, register access, interrupts, DMA, etc.
- CXL.cache: Device access to processor memory
- CXL.memory: Processor access to device attached memory

* CXL Type1&2 device support dependent on ecosystem readiness; type 2 PoC only

Процессоры Intel Xeon 6

Процессоры Intel Xeon 6

Intel® Xeon® 6 Processors

Launching
6/3 8pm PDT - 6/4 11am CST

Q3'2024

Q1'2025



Intel® Xeon® 6700E



Intel® Xeon® 6900P



Intel® Xeon® 6900E

Intel® Xeon® 6700P
Intel® Xeon® 6500P
Intel® Xeon® 6 SoC
Intel® Xeon® 6300P

Процессоры Intel Xeon 6

Four categories of workload-optimized processors

Intel Xeon 6700-series with E-cores

- Launching now
- Up to 144 cores (144 threads) per CPU
- Up to 330W per CPU

Intel Xeon 6900-series with E-cores

- Future release
- Up to 288 cores (288 threads) per CPU
- Up to 500W per CPU

↑
Pin Compatible
(LGA4710)
↓

↑
Pin Compatible
(LGA7529)
↓

Intel Xeon 6700-series with P-cores

- Future release
- Up to 86 cores (172 threads) per CPU
- Up to 350W per CPU

Intel Xeon 6900-series with P-cores

- Future release
- Up to 128 cores (256 threads) per CPU
- Up to 500W per CPU

Процессоры Intel Xeon 6

INTEL® XEON® 6 PROCESSORS Performance Core




INTEL® XEON® 6 PROCESSORS Efficient Core

up to 86 cores (6700) or 128 cores (6900)	Cores	up to 144 cores (6700) or 288 cores (6900)
1S, 2S, 4S, 8S	Sockets	1S, 2S
125 to 500W	Max TDP	205 to 500W
up to 12 channels DDR5 MCR DIMM	Memory	up to 12 channels DDR5
6400 (1 DPC) 5200 (2 DPC) 8800 MCR (1 DPC)	Max Memory Speed	6400 (1 DPC) 5200 (2 DPC)
up to 6 UPI 2.0 up to 24 GT/s per lane	Intel® UPI	up to 4 UPI 2.0 up to 24 GT/s per lane
up to 96 lanes PCIe 5.0 (x16, x8, x4, x2) R1S: up to 136 lanes PCIe 5.0 for single socket designs	PCI Express	up to 96 lanes PCIe 5.0 (x16, x8, x4, x2)
up to 64 lanes CXL 2.0	Compute Express Link	up to 64 lanes CXL 2.0
52/57	Physical/Virtual Address Bits	52/48
Intel Advanced Matrix Extensions (INT8, BF16, FP16) Intel Advanced Vector Extensions 512 (VNNI/INT8)	AI Acceleration Intel® Deep Learning Boost	Intel Advanced Vector Extensions 2 (VNNI/INT8)
Intel Software Guard Extensions, Intel Trusted Domain Extensions	Security	Intel Software Guard Extensions, Intel Trusted Domain Extensions
Vector AES, SHA2-256 extensions, VPMADD52	Crypto	Vector AES, SHA2-256 extensions, VPMADD52
Intel QuickAssist Technology, Intel Dynamic Load Balancer , Intel Data Streaming Accelerator , Intel In-memory Analytics Accelerator	Integrated Accelerators	Intel QuickAssist Technology, Intel Dynamic Load Balancer , Intel Data Streaming Accelerator , Intel In-memory Analytics Accelerator

Процессоры Intel Xeon 6

Addressing Unique Workload Requirements



Процессоры Intel Xeon 6 (6700-series)

Intel® Xeon® 6 E-core SKU Map

SKU	Cores	Micro Architecture	Base (GHz)	All Core Turbo (GHz)	Max Turbo (GHz)	L3 Cache (MB)	TDP (Watts)	Max Scala.	DDR5 Memory Speed (1DPC)	Default Accelerator Devices	Intel TDX Keys (Per CPU)	Long Life Available*	UPI Links Enab.	PCIe5 Express Lanes/CXL
6780E	144	E-core	2.2	3.0	3.0	108	330	2S	6400	2 Intel® DSA, 2 Intel® IAA, 2 Intel® QAT, 2 Intel® DLB	2048	✓	4	88
6766E	144	E-core	1.9	2.7	2.7	108	250	2S	6400	2 Intel DSA, 2 Intel IAA, 2 Intel QAT, 2 Intel DLB	1024	✓	4	88
6756E	128	E-core	1.8	2.6	2.6	96	225	2S	6400	2 Intel DSA, 2 Intel IAA, 2 Intel QAT, 2 Intel DLB	1024	-	4	88
6746E	112	E-core	2.0	2.7	2.7	96	250	2S	5600	2 Intel DSA, 2 Intel IAA, 2 Intel QAT, 2 Intel DLB	1024	-	4	88
6740E	96	E-core	2.4	3.2	3.2	96	250	2S	6400	2 Intel DSA, 2 Intel IAA, 4 Intel QAT, 4 Intel DLB	1024	✓	4	88
6731E	96	E-core	2.2	3.1	3.1	96	250	1S	5600	2 Intel DSA, 2 Intel IAA, 2 Intel QAT, 2 Intel DLB	1024	-	0	88
6710E	64	E-core	2.4	3.2	3.2	96	205	2S	5600	2 Intel DSA, 2 Intel IAA, 4 Intel QAT, 4 Intel DLB	1024	✓	4	88

*Long Life Availability: 7+ years



Intel may make changes to specifications and product descriptions at any time, without notice.

Please visit [intel.com/xeon](https://www.intel.com/xeon) or contact your Intel representative to obtain the latest product specifications. Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. All processors support Intel Virtualization Technology (Intel VT-x).

Intel Confidential

Процессоры Intel Xeon 6 (6900-series)



Процессоры Intel Xeon 6 (6900-series)

Intel Xeon 6 P-core SKU Map

Roadmap SKUs with customization options

PERFORMANCE SKUs


SKU	CORES	BASE (GHz)	ALL CORE TURBO (GHz)	Max TURBO (GHz)	L3 CACHE (MB)	TDP (Watts)	Max. Scala.	Memory Channels	DDR5 Memory Speed	MRDIMM Speed	Default Accel. Devices	Intel TDX Keys (Per CPU)	UPI Links Enab.	PCIe Lanes
6980P	128	2.0	3.2	3.9	504	500	2S	12	6400	8800	4/4/4/4	1024	6	96
6979P	120	2.1	3.2	3.9	504	500								
6972P	96	2.4	3.5	3.9	480	500								
6952P	96	2.1	3.2	3.9	480	400								
6960P	72	2.7	3.8	3.9	432	500								

Intel may make changes to specifications and product descriptions at any time, without notice.


Please visit [intel.com/xeon](https://www.intel.com/xeon) or contact your Intel representative to obtain the latest product specifications. Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. All processors support Intel Virtualization Technology (Intel VT-x).

*Accelerators List Order:
DSA, IAA, QAT, DLB
Intel® AMX featured in each core

Процессоры Intel Xeon 6 (6900-series)



Процессоры Intel Xeon 6 (6900-series)




Процессоры Intel Xeon 6 (6900-series)

Multiplexed Rank DIMMs

First to market on Intel Xeon 6 processors with P-core



Intel® Xeon® 6 with P-cores (128c)
MRDIMM-8800 Performance Gains Over DDR5-6400
Higher is better



Процессоры Intel Xeon 6 (6900-series) память MRDIMM



Память DDR5-6400

Память MRDIMM-8800

Сравнительные бенчмарки процессоров AMD Turin и Intel Xeon 6


Молекулярная динамика

GROMACS 2024

Implementation: MPI CPU - Input: water_GMX50_bare




Phoronix.com



Computational Fluid Dynamics (CFD) – небольшая сетка

OpenFOAM 10

Input: driveerFastback, Small Mesh Size - Execution Time




Computational Fluid Dynamics (CFD) – средняя сетка

OpenFOAM 10

Input: drivaeFastback, Medium Mesh Size - Execution Time



Phoronix.com




HPCG benchmark – набор высокопроизводительных тестов

High Performance Conjugate Gradient 3.1

X Y Z: 144 144 144 - RT: 60

pts

Phoronix.com




Типы CPU-систем под новые процессоры

Классические серверы в стойку

1RU = 1.75" = 44 мм

↓
1 RU
4 x 3.5" disks



↓
1 RU
10+ x 2.5" disks



↓
1 RU
12 x 3.5" disks



↓
2 RU
24 x 2.5" disks



ASUS →




GigaByte →




Gooxi →

Универсальный 2RU сервер с 12 x 2.5"/3.5" дисками




2 x PSU 1000W+
5+ x PCIe 4.0 (5.0)

Типичная блок-схема сервера



Типичная блок-схема сервера (упрощенная)




Поколения процессоров и их важные отличия

AMD EPYC							
<i>Generation</i>	<i>Naming</i>	<i>Max P-cores</i>	<i>Max E-cores</i>	<i>DRAM Gen</i>	<i>DRAM Channels</i>	<i>PCIe Gen</i>	<i>PCIe lanes</i>
Zen 3 (Milan)	7xx3	64		DDR4	8	Gen 4.0	128
Zen 4 (Genoa)	9xx4	96	128	DDR5	12	Gen 5.0	128
Zen 5 (Turin)	9xx5	128	192	DDR5	12	Gen 5.0	160
Intel							
<i>Generation</i>	<i>Naming</i>	<i>Max P-cores</i>	<i>Max E-cores</i>	<i>DRAM Gen</i>	<i>DRAM Channels</i>	<i>PCIe Gen</i>	<i>PCIe lanes</i>
3th	x3xx	40		DDR4	8	Gen 4.0	64
4th-5th	x4xx/x5xxx	60/64		DDR5	8	Gen 5.0	80
6th	67xx/69xx	128	144	DDR5	12	Gen 5.0	88/96

Дисковое хранение данных (СХД)

Современные топовые SSD диски для СХД



Современные топовые SSD диски для СХД



Современный топовые SSD диски для СХД

Solidigm™ D5-P5336 overview



Features

Product Name	Solidigm D5-P5336			
Media	192L QLC NAND (171GiB)			
Power off Retention	3 months @ 40°C			
Indirection Unit	16KB			
User Capacity	7.68TB	15.36TB	30.72TB	61.44TB
Endurance (5-yr DWPD) ¹	0.42	0.51	0.56	0.58
Endurance (PBW) ¹	5.9	14.1	31.5	65.2
Max Power	25 W			
Idle Power	<5 W			
UBER	< 1 Sector per 10 ¹⁷ Bits Read			
MTBF	2 Million Hours			
Features	OCP 2.0 support ² , NVMe 1.4 Compliance ³ , FIPS 140-3 Level 2			


Performance

4K Random Read, IOPS, QD256	up to 1.005M
16K Random Write, IOPS, QD256	up to 43K
128K Seq. Read, MB/s, QD128	up to 7,000
128K Seq. Write, MB/s, QD128	up to 3,300

1. IU-Aligned Endurance. Based on 100% Random Write 16KB workload calculation.

2. See Solidigm D5-P5336 product specification for exceptions and modifications for compliance/support details.

3. PRQ1 SKUs compliant with NVMe version 1.4 and NVMe MI 1.1. NVMe 2.0 and NVMe MI 1.2 support to be targeted in PRQ2 and subsequent releases.



U.2 (15mm)



E3.S (7.5mm)



E1.L (9.5mm)



New Spec Alignment

NVMe 1.4c³ (Target NVMe 2.0 support for PRQ2) and OCP 2.0 Support² (Latency monitor, FW history log, NSSR, Format Progress Indicator, NUSE specific etc.)



PCIe 4.0 Controller


Delivers better latency, expanded management capabilities, and critical new NVMe features compared to previous gen QLC SSD.



Data Center SSD features

New Trim arch & improvements, new OCP 2.0 feature support, SGL, VSS, HDR with Opal, SGL with DSM, Opal Locking Range with MNS, Extended DST, FIPS 140-3 Level 2 (Future certification on generic SKUs), Telemetry, etc.

Современные топовые SSD диски для СХД



Современные топовые SSD диски для СХД



Диск Micron 6550 ION:

слева - U2;
внизу- E3S.



Современные топовые SSD диски для СХД

Скоростные характеристики диска Micron 6550 ION

World's fastest 60TB data center SSD

Best-in-class performance versus 60TB competition using up to 20% less power (20W vs. 25W)

**12.0
GB/s**

sequential reads
up to 179% better

**5.0
GB/s**

sequential writes
up to 150% better


**1,600
KIOPS**

random reads
up to 80% faster


**3.4
hours**

competition takes
up to 150% longer

Sequential read throughput (128KB)



Sequential write throughput (128KB)




Random read throughput (4KB)



Fill rate

(100% 128KB sequential writes)



1. The Micron 6550 ION offers a capacity of up to 61.44TB. Comparisons are made with other 61.44TB NVMe SSDs from Samsung, Solidigm, and Western Digital (WD). These comparisons use publicly available competitor information from published sources at the time of the Micron 6550 ION launch.
2. The Micron 6550 ION 61.44TB SSD at 20W provides better performance/watt than 61.44TB SSD's at 25W from Solidigm and Samsung and 20W from Western Digital, resulting in up to 20% power savings, respectively. Performance per watt claims are vs. competitive drives data sheet specifications per footnote 1.
3. Fill time calculation = 100% 128KB sequential writes at spec noted on public datasheet until drive's usable capacity is full.

Современные топовые HDD диски для СХД



Шпиндельные диски 3.5:

32ТБ - WD Ultrastar DC HC690

36ТБ - Seagate Mozaic 3+



Система хранения на 60 дисков 3.5"




SSG-640SP-E1CR60

Последние достижения сетевых технологий

Infiniband

(Backend network)

InfiniBand Roadmap



*Link speeds specified in Gb/s at 4X (4 lanes)

Коммутатор NDR IB на 64 порта x 400Gbit/s

ANNOUNCING NVIDIA QUANTUM-2

400G NDR InfiniBand
Cloud-Native Supercomputing



QUANTUM-2 SWITCH

57 Billion Transistors TSMC 7N

Optimized Multi-Tenant In-Network Computing

64-Ports of 400 Gbps or 128-Ports of 200 Gbps

3X Higher Switching Throughput | 32X More AI Acceleration Engines

Sampling Now



Сетевая карта NDR и разновидности кабелей



Подключение



Собранный кабель 400G - 400G




Оптический патч-корд 400G - 400G



Оптический патч-корд 400G - 4 x100G

Трансиверы



QSFP-DD
8-Channel
12W

QSFP28
4-Channel
3.5W

SFP-DD
2-Channel
3.5W

Собранный кабель 400G - 4 x100G

Не-nVidia сообщество в качестве Backend-сети предлагает UltraEthernet

UltraEthernet
Consortium

Evolve ethernet as an open, interoperable, high performance,
full-communications stack architecture to meet the growing
network demands of AI and HPC at scale

UEC 1.0 Specification - Q1CY25

Performant

Scalable

Cost Effective

Сообщество и оно же - поддержка UltraEthernet

UltraEthernet
Consortium

Steering Members

AMD

ARISTA

BROADCOM

CISCO

EVIDEN
an atos business

Hewlett Packard
Enterprise

intel

Meta

Microsoft

ORACLE

Ethernet

(Frontend network)

Презентация экосистемы 400GbE nVidia в 2022 г

ANNOUNCING NVIDIA SPECTRUM ETHERNET PLATFORM



END-TO-END 400G HYPERSCALE NETWORKING PLATFORM

Performance Optimized for Cloud, Enterprise, AI & Simulation at Scale


SPECTRUM-4 SWITCH



Accelerated Ethernet Cloud Fabric
Cloud-Native, AI & Accelerated Computing

Up to 800Gbs

ConnectX-7 SMARTNIC



The World's Leading Network Adapter
Intelligent Offloads · RDMA · Precision Timing

Up to 400Gbs

BLUEFIELD-3 DPU



Programmable Data Center Infrastructure
Accelerated Networking, Security, Storage

Up to 400Gbs

Презентация экосистемы 400GbE nVidia в 2022 г

ANNOUNCING SPECTRUM-4 ETHERNET SWITCH

Accelerated Ethernet Fabric — Cloud, Enterprise, and AI



51.2T
Bandwidth

100B
Transistors


12.8T
Crypto

37.6B
Packet Rate


**400G &
800G**
Ports

4X
Higher
Throughput


Визуализация соединений 400GbE в ЦОДе




Рoadmap Ethernet-коммутаторов nVidia Spectrum (2024г)



Чип 400GbE-коммутатора Tomahawk 4 компании Broadcom



64 x 400GbE - коммутатор FS-N9510-64D на чипе Tomahawk 4



Рoadmap Ethernet-коммутаторов на чипах Broadcom Tomahawk

Broadcom Ethernet Leadership


Long-term Roadmap with Proven Execution

100x Bandwidth Increase

- Outpacing Moore's Law

14x Power/Performance Increase


- Consistently lowest Watts/Gbps



Рoadmap Ethernet на примере оптических трансиверов

ARISTA

Pluggable Optics Modules Roadmap



Системы с ускорителями GPU



Accelerate Everything

GPU Optimized Systems to Achieve 5X, 10X,... 100X Performance



Large Scale AI Training Workloads

Large language models, Generative AI training, autonomous driving, robotics



HPC/AI Workloads

Engineering simulation, scientific research, genomic sequencing, drug discovery



Enterprise AI Inference & Training

AI-enabled services/applications, chatbots, business automation



Visualization and Design

Graphical content development and automatic generation, digital twins, 3D collaboration



Content Delivery and Virtualization

Content delivery networks (CDNs), video transcoding, live streaming, VDI




AI Edge

Retail automation, manufacturing/logistics automation, medical diagnosis/predictive care, security, and many more

AI - One Architecture | Train and Deploy Everywhere

One-Year Rhythm



GPU – разные производители похожие форм-факторы



Multi socket


HGX

H100
H200
B100
B200

Coming
Soon



CDNA3
MI300X



Gaudi3 UBB



PCIe

H100 NVL
L40S
L4



MI210



Gaudi3 PCIe



CPU+GPU


Grace Hooper

GH200
GB200


Coming
Soon



CDNA3
MI300A



Разные GPU для разных рабочих нагрузок

Manufacturer	GPU Model	Architecture	 DL Training & DA	 DL Inference	 HPC / AI	 Omniverse / Render Farms	 AI Video	 Far Edge Acceleration
 NVIDIA.	H200	Multi Socket	●	●	●			
 NVIDIA.	H100	Multi Socket	●	●	●			
 AMD	MI300X	Multi Socket	●	●				
 intel.	GAUDI3	Multi Socket	●	●	●			
 NVIDIA.	H100NVL	PCIe	●	●				
 NVIDIA.	L40S	PCIe	●	●	●	●	●	
 AMD	MI300A	CPU+GPU			●			
 intel.	GAUDI3	PCIe		●	●			
 NVIDIA.	L4	PCIe		●		●	●	●
 NVIDIA.	GH200	CPU+GPU	●	●	●			



Best



Better



Good

Price-performance comparison relative across each entire workload column. This chart should be used in conjunction with measured data for targeted workloads.

Сравнение производительностей GPU nVidia разных поколений

Specification	NVIDIA B200	NVIDIA B100	NVIDIA H200	NVIDIA H100	NVIDIA A100
Architecture	Blackwell	Blackwell	Hopper	Hopper	Ampere
GPU Name	NVIDIA B200	NVIDIA B100	NVIDIA H200	NVIDIA H100	NVIDIA A100
FP64	40 teraFLOPS	30 teraFLOPS	34 teraFLOPS	34 teraFLOPS	9.7 teraFLOPS
FP64 Tensor Core	40 teraFLOPS	30 teraFLOPS	67 teraFLOPS	67 teraFLOPS	19.5 teraFLOPS
FP32	80 teraFLOPS	60 teraFLOPS	67 teraFLOPS	67 teraFLOPS	19.5 teraFLOPS
FP32 Tensor Core	2.2 petaFLOPS	1.8 petaFLOPS	989 teraFLOPS	989 teraFLOPS	312 teraFLOPS
FP16/BF16 Tensor Core	4.5 petaFLOPS	3.5 petaFLOPS	1979 teraFLOPS	1979 teraFLOPS	624 teraFLOPS
INT8 Tensor Core	9 petaOPs	7 petaOPs	3958 teraOPs	3958 teraOPs	1248 teraOPs
FP8 Tensor Core	9 petaFLOPS	7 petaFLOPS	3958 teraFLOPS	3958 teraFLOPS	-
FP4 Tensor Core	18 petaFLOPS	14 petaFLOPS	-	-	-
GPU Memory	192GB HBM3e	192GB HBM3e	141GB HBM3e	80GB HBM3	80GB HBM2e
Memory Bandwidth	8TB/s	8TB/s	4.8TB/s	3.2TB/s	2TB/s


GPU-ускорители форм-фактора PCIe

GPU-ускорители форм-фактора PCIe




GPU-ускорители максимального класса

- nVidia H100 80GB PCIe
- nVidia H100 94GB PCIe NVL
- nVidia H200 141GB PCIe NVL




Две карты H100 объединены в одну с помощью внешнего соединения Nvlink

PCIe GPU ускорители nVidia среднего класса



nVidia L40S 48GB
(пассивное охлаждение)

nVidia RTX 6000 Ada 48GB
nVidia RTX Pro 6000 96GB
(активное охлаждение)



GPU-ускорители среднего класса форм-фактора PCIe

NVIDIA Ada Lovelace Architecture

	<u>RTX 6000</u>	<u>RTX 5000</u>	<u>RTX 4500</u>	<u>RTX 4000</u>
GPU Memory	48GB GDDR6 with ECC	32GB GDDR6 with ECC	24GB GDDR6 with ECC	20GB GDDR6 with ECC
Display Ports	4x DisplayPort 1.4a*	4x DisplayPort 1.4a*	4x DisplayPort 1.4a	4x DisplayPort 1.4a
Max Power Consumption	300W	250W	210W	130W
Form Factor	4.4" (H) x 10.5" (L) dual slot	4.4" (H) x 10.5" (L) dual slot	4.4" (H) x 10.5" (L) dual slot	4.4" (H) x 9.5" (L) single slot
Thermal	Active	Active	Active	Active
	Learn More >	Learn More >	Learn More >	Learn More >

Новые GPU-ускорители среднего класса форм-фактора PCIe

NVIDIA Blackwell Architecture

	<u>RTX PRO 6000 Blackwell</u> <u>Workstation Edition</u>	<u>RTX PRO 6000 Blackwell</u> <u>Max-Q Workstation</u> <u>Edition</u>	<u>RTX PRO 5000</u> <u>Blackwell</u>	<u>RTX PRO 4500</u> <u>Blackwell</u>
GPU Memory	96GB GDDR7 with ECC	96GB GDDR7 with ECC	48GB GDDR7 with ECC	32GB GDDR7 with ECC
Display Ports	4x DisplayPort 2.1	4x DisplayPort 2.1	4x DisplayPort 2.1	4x DisplayPort 2.1
Max Power Consumption	600W	300W	300W	200W
Form Factor	5.4" (H) x 12.0" (L) dual slot	4.4" (H) x 10.5" (L) dual slot	4.4" (H) x 10.5" (L) dual slot	4.4" (H) x 10.5" (L) dual slot
Thermal	Double Flow Through	Active	Active	Active
	Learn More >	Learn More >	Learn More >	Learn More >

Сравнение GPU-карт в форм-факторе PCIe

Тип карты	Память (GB)	FP64 (Tflops)	FP32(Tflops)	FP16 (Tflops)	Охлаждение	Цена за FP32 (ИИ)
<i>Data-center Top</i>						
A100 PCIe	80	10	156	78	Пассивное	246
H100 PCIe	80	30	837	1,671	Пассивное	54
H100 NVL PCIe	94	30	837	1,671	Пассивное	54
H200 NVL PCIe	141	30	837	1,671	Пассивное	64
<i>Data-center Pro</i>						
L40S	48	1.4	91	91	Пассивное	160
<i>RTX Pro (Blackwell)</i>						
Pro 6000	96	2.0	126	126	Активное	151
Pro 5000	48	1.1	74	74	Активное	125
Pro 4000	24	0.7	47	47	Активное	67
<i>RTX Ada (Ada Lovelace)</i>						
Ada 6000	48	1.4	91	91	Активное	140
Ada 5000	32	1.0	65	65	Активное	115
Ada 4500	24	0.6	40	40	Активное	83
<i>Desktop</i>						
RTX 5080 Super	24	0.9	56	56	Активное	41
RTX 4070 Ti Super	16	0.7	44	44	Активное	36

GPU-ускорители форм-фактора SXM (Nvlink) и OAM

GPU – разные производители похожие форм-факторы



Multi socket


HGX

H100
H200
B100
B200

Coming
Soon



CDNA3
MI300X



Gaudi3 UBB



PCIe

H100 NVL
L40S
L4



MI210



Gaudi3 PCIe



CPU+GPU


Grace Hooper

GH200
GB200

Coming
Soon



CDNA3
MI300A



GPU-система 8U на 8 карт GPU H100/H200 SXM (Supermicro)


(Angled View – System)



16x 2.5" Hot-swap NVMe Drive Bays + 3x 2.5" Hot-swap SATA Drive Bays

На процессорах:
Intel -> SYS-821GE-TNHR
AMD -> AS-8125GS-TNHR

(Rear View – System)



	Slot Description
1	PCIe 5.0 x16 (LP) from PLX switch linked to GPUs
2	PCIe 5.0 x16 (LP) from PLX switch linked to GPUs
3	PCIe 5.0 x16 (LP) from PLX switch linked to GPUs
4	PCIe 5.0 x16 (LP) from PLX switch linked to GPUs
5	PCIe 5.0 x16 (LP) from PLX switch linked to GPUs
6	PCIe 5.0 x16 (LP) from PLX switch linked to GPUs

	Slot Description
7	PCIe 5.0 x16 (LP) from PLX switch linked to GPUs
8	PCIe 5.0 x16 (LP) from PLX switch linked to GPUs
9	PCIe 5.0 x16 (FHHL) via PLX switch (optional)
10	PCIe 5.0 x16 (FHHL) via PLX switch (optional)
11	PCIe 5.0 x16 (FHHL) via PLX switch
12	PCIe 5.0 x16 (FHHL) via PLX switch


GPU-системы 8U на 8 карт GPU H100/H200 SXM (8xGPU HGX – модуль выдвинут)



Gigabyte
G593-SD0

Система 8 x B200 (180GB) на процессорах Intel Xeon 6900

SYS-A22GA-NBRT




10 Hot-swap
Fan Modules

6 Redundant 5000W Titanium
Level Power Supplies


2 10G RJ45 Ports
Dedicated IPMI Port

GPU-системы 8U на 8 карт GPU H100/H200 SMX

(Top View – System)



(Top View – System)



AS -8125GS-TNMR2

MI300X Solution



Specifications

CPU –

Dual EPYC 9004 Series processors, up to 256 cores/512 threads

Memory –

24x DIMM slots, ECC DDR5 supports up to 4800MHz

Drives –

2x onboard NVMe M.2
12x PCIe5.0 x4 NVMe U.2
(additional optional) 4x PCIe5.0 x4 NVMe U.2
2x SATA 2.5"

Expansion –

8x PCIe 5.0 x16 low profile
2x FHFL PCIe 5.0 x16 slots
(Optional) 2x PCIe 5.0 x16 slots via additional PCIe switch

System Cooling

5x Front and 5x Rear counter-rotating fans with optimal fan speed control

Power Supply –


6x or 8x 3000-watt N+N Redundant Titanium Level Power Supplies

Accelerator Support -

8x MI300X

Dimensions – W x H x D

14.0" (H) x 17.2" (W) x 33.2" (D)



Recommend System Configuration:

Item#	Description	Q'ty
AS -8125GS-TNMR2	[NR]H13DSG-OM, CSE-GP801TS-D2 for MI300X	1
PSE-GEN9534-0799	Genoa 9534 DP/UP 64C/128T 2.45G 256M 280W SP5	2
MEM-DR564L-CL01-ER48	64GB DDR5-4800 2RX4 (10X4) LP (16Gb) ECC RDIMM	24
HDS-MMN-MTFDKBA960TFR-15	Micron 7450 PRO 960GB NVMe PCIe 4.0 M.2 22x80mm TCG Opal	2
AOC-CX766003N-SQ0	Nvidia 900-9X766-003N-SQ0 PCIe 1-port IB 400GE OSFP Gen5	8
GPU-AMD-MI300X-OAM-0045H	[NR] AMD Instinct MI300X 192GB 8 OAM + UBB	1
AOC-STGS-I2T-O	Std LP 2-port 10G RJ45, Intel X550 (Retail Pack)	1

SMCI Gaudi3 Product Spec

Product Model : SYS-822GA-NGR3

Key Features

- Supports 8 Gaudi3 HL-325 OAM with 128GB HBM
- Direct connect on-board 800GbE OSFP-DD Scale-out
- Dual 6rd Gen Series Intel Xeon Scalable Processors

Key Application

- AI Compute/Model Training/Deep Learning

Specification

CPU – Dual Socket

Dual P-Cores Granite Rapids AP and E-Cores
Clearwater Forest-AP processor up to 500W

Memory – 24 DIMM Slots

24x DIMM ECC DDR5 designed for up to 6400
MT/s (1DPC)

Drives – 8 Hot-Swap Bays

8x NVMe Gen 5
2x M.2 2280/22110

Expansion – 5 PCI-E Gen5 Slots

2xPCI-E 5.0 x16 (FHHL)
2xPCI-E 5.0 x 8 (FHHL)
1xAIOM with 10G LAN card (2x Intel i550)

High-Speed QSFP-DD– 800GbE

6x OSFP-DD ports
1x RJ45 1GbE IPMI

Power Supply – 4+4 Redundant Power


3000W Titanium Level

Intel Gaudi3
HL-325L 8-HPU

6x 800GbE OSFP-DD
Scale-Out

PCIE
4x PCIE Gen5 slot


2.5" Drive Bays
8x NVMe Hot-Swap



Front Side

Redundant
FAM Models


Redundant
Power Supplies



Rear Side

Система 8 x Intel Gaudi3 на процессорах Intel Xeon 6900

SYS-822GA-NGR3



(nVidia GTC-2025)

Улетаем в будущее

– новые ускорители:

- Blackwell**

- Rubin**


- Feinman**

- новые профессиональные «десктопы»

- DGX Spark**

- DGX Station**

Идея Дженсена Хуанга – «все есть GenerativeAI»



«Дата-центры – это фабрики по производству токенов»



Масштабируем GPU-системы, чтобы фабрики токенов были эффективными



NVIDIA EOS

18 DGX PODS

FP8	18 EFLOPS	6X
FP16	9 EFLOPS	3X
FP64	275 PFLOPS	3X
In-Network Compute	3.7 PFLOPS	36X
Bisection Bandwidth	230 TB/s	2X

Оценка стоимости обучения Chat GPT 4.0 1.8T

GPT-4 1.8T MoE	2016	2018	2020	2022	2024
10 Day Training	"Pascal"	"Volta"	"Ampere"	"Hopper"	"Blackwell"
GPU	P100	V100	A100	H100	B100
Peak Teraflops	19	130	620	4,000	20,000
<u>Training And Inference Precision</u>	<u>FP16</u>	<u>FP16</u>	<u>FP16</u>	<u>FP8</u>	<u>FP4</u>
Inference Joules/Token	17,000	1,200	150	10	0.4
GPUs to Train In 10 Days	42,105,300	6,153,800	1,290,300	100,000	10,000
GPU Unit Price	\$5,000	\$10,000	\$15,000	\$27,500	\$37,500
<u>GPU Cost For Training In 10 Days</u>	<u>\$210.5 B</u>	<u>\$61.5 B</u>	<u>\$19.4 B</u>	<u>\$2.8 B</u>	<u>\$0.4 B</u>
Power To Train, Gw-hour	1,000	140	40	13	3
Electricity Cost Per Kw-hour	\$0.14	\$0.14	\$0.14	\$0.16	\$0.18
<u>Electricity Cost For Training Run</u>	<u>\$140.0 M</u>	<u>\$19.6 M</u>	<u>\$5.6 M</u>	<u>\$2.1 M</u>	<u>\$0.54 M</u>
<u>Three Year Electricity Cost</u>	<u>\$15.3 B</u>	<u>\$2.1 B</u>	<u>\$613.6 M</u>	<u>\$227.9 M</u>	<u>\$59.2 M</u>

Blackwell – форм-факторы GPU



GB200 NVL72

Compute for Trillion Parameter Scale AI
Maximum Performance and Lowest TCO



HGX B200

Best Performance and TCO for HGX Platform



HGX B100

Drop-in Upgrade for Existing Hopper Infrastructure

GB200 NVL72 Compute and Interconnect Nodes

Building Blocks for the GB200 NVL72 Rack



GB200 SUPERCHIP

40 PETAFLOPS FP4 AI INFERENCE
20 PETAFLOPS FP8 AI TRAINING
864GB FAST MEMORY



GB200 SUPERCHIP COMPUTE TRAY

2x GB200
80 PETAFLOPS FP4 AI INFERENCE
40 PETAFLOPS FP8 AI TRAINING
1728 GB FAST MEMORY
1U Liquid Cooled
18 Per Rack



NVLINK SWITCH TRAY

2x NVLINK SWITCH CHIP
14.4 TB/s Total Bandwidth
SHARPV4 FP64/32/16/8
1U Liquid Cooled
9 Per Rack


Blackwell Ultra NVL72 – единая система (больше GPU) из 72 GPU Blackwell Ultra и 36 CPU Grace



Blackwell Ultra NVL72
Second Half 2025

- 1.1 EF Dense FP4 Inference
- 0.36 EF FP8 Training
- 1.5X GB200 NVL72
- New Attention Instructions
2X
- 20 TB HBM | 40 TB Fast Memory
1.5X
- 14.4 TB/s CX8
2X

Grace




Blackwell Ultra



2 Reticle-Sized GPUs
1.1 EF Dense FP4 | 288GB HBM3e

Oberon Rack
Liquid Cooled

Vera Rubin NVL144 – единая система из 144 чипов Rubin и CPU на 36 чипах Vera



Vera Rubin NVL144
Second Half 2026

- 3.6 EF FP4 Inference
- 1.2 EF FP8 Training
- 3.3X GB300 NVL72
- 13 TB/s HBM4
- 75 TB Fast Memory
- 1.6X
- 260 TB/s NVLink6
- 2X
- 28.8 TB/s CX9
- 2X

Vera



88 Custom Arm Cores
176 Threads
1.8 TB/s NVLink-C2C

Rubin



Die-Sized GPUs
FP4 | 288GB HBM4

Oberon Rack
Liquid Cooled



Rubin Ultra NVL576 – единая система из 576 чипов Rubin Ultra и CPU на 144 чипах Vera



Rubin Ultra NVL576
Second Half 2027

- 15 EF FP4 Inference
- 5 EF FP8 Training
- 14X GB300 NVL72
- 4.6 PB/s HBM4e
- 365 TB Fast Memory
- 8X
- 1.5 PB/s NVLink7
- 12X
- 115.2 TB/s CX9
- 8X

*Kyber Rack
Liquid Cooled*

Vera



- 88 Custom Arm Cores
- 176 Threads
- 1.8 TB/s NVLink-C2C


Rubin Ultra




1e-Sized GPUs
FP4 | 1TB HBM4e




Рoadmap GPU nVidia– появился чип Feynman (2028)



Новые вычислители потребуют новой сети: кремниевые-фотонные коммутаторы





Настольный ускоритель - nVidia DGX Spark




- > NVIDIA GB10 Grace Blackwell superchip
- > 1,000 AI TOPS of FP4 AI performance
- > 128GB of coherent, unified system memory
- > ConnectX-7 Smart NIC
- > Up to 4TB Storage
- > 150mm L x 150mm W x 50.5mm H

Настольный ускоритель - nVidia DGX Spark



Рабочая станция - nVidia DGX Station



NVIDIA GPU	1x NVIDIA Blackwell Ultra
NVIDIA CPU	1x Grace-72 Core Neoverse V2
GPU Memory	Up to 288GB HBM3e 8 TB/s
CPU Memory	Up to 496GB LPDDR5X Up to 396 GB/s
NVLink-C2C	Up to 900 GB/s
Networking Peak Bandwidth	NVIDIA ConnectX®-8 SuperNIC Up to 800 Gb/s

FIN