ZEUS:

A New GPU for High Performance Workloads

March 6, 2025





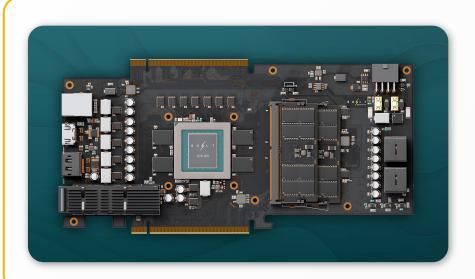


INTRODUCTION



ZEUS

The Only GPU Designed for High Performance Workloads



Our customers came to us asking for a better 3D graphics experience and to reduce their power consumption.

So we completely re-designed the GPU to give architects, gamers, artists, and researchers what they need to change the world.

What Our Customers Do



and why they asked for a better GPU:

Rendering



Artists, designers, and engineers rely on accurate renderings to explore and share ideas.

Zeus reduces the time they spend waiting for renders.

Real Time Path Tracing requires: 28 Zeus GPUs vs 280 legacy GPUs

Gaming



Gamers demand higher frame rates, pushing devs to work around performance limitations of today's GPUs.

Developers spend less time optimizing and gamers reap the benefits.

HPC



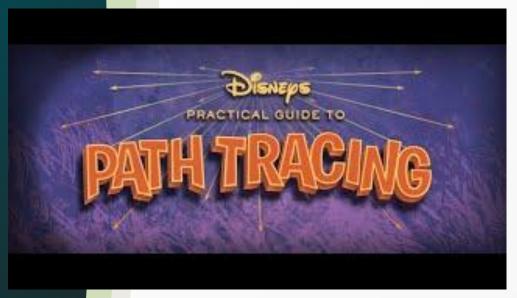
Engineers and researchers rely on precise simulations to research and design parts and products.

Zeus expands simulation possibilities while reducing power consumption.

44 KW vs. 130 KW legacy GPUs per rack



Path Tracing 101



Disney's Practical Guide to Path Tracing

- Path tracing is an advanced form of ray tracing that more accurately simulates light
- More computationally expensive vs ray tracing
- Lots of rays required to converge on clean results (reduce noise)
- Legacy GPUs don't have the performance needed to compute enough rays for real time path tracing

Upgraded Working View









Working View with legacy GPUs

Working View with Zeus

Final Render

Legacy GPUs cannot path trace in real time without major workarounds.

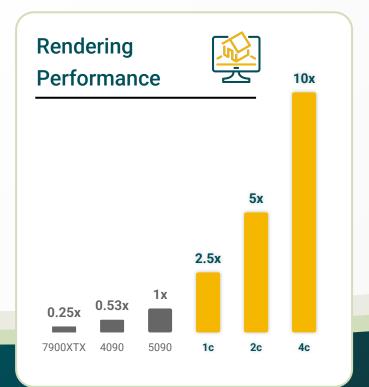
The image the creators see in their working view does not look like the final render. It leaves out critical visual information like lighting effects, fur, smoke, fire, physical materials, and more.



RENDERING

Rendering Performance Improvements





Improved rendering performance opens up new workflows of truly collaborative creation with clients, coworkers, and audiences. Creators can iterate through ideas in real time with their team instead of waiting hours to see what they've changed. Zeus gives creators the time they need to make adjustments without missing a deadline.





Benchmark Notes



- Pre-silicon performance numbers
- Rendering performance: Measured as the time to path trace a scene
 - Sponza (<u>source</u>) with curtains & ivy addons @ 1080p 100spp
 - Imported as USD into renderer same geometry + materials when possible
 - Nvidia Bench: AMD Ryzen 9 5950X, 128 GB DDR4-2133, RTX 4090 (566.36, no OC), Windows
 11 Pro 22631.4751, Omniverse USD Composer 2023.2.5 (Iray)
 - AMD Bench: AMD Ryzen 9 5950X, 128 GB DDR4-2133, 7900 XTX (24.12.1, no OC), Windows
 11 Pro 22631.4751, Blender 4.2.2. LTS Cycles (HIP enabled)
 - Rendered without materials as Cycles has limited support for MaterialX
 - Bolt Bench: AMD Ryzen 9 7900X, 128 GB DDR5-3600, Xilinx U50 FPGA, Ubuntu 22.04.2 LTS,
 Bolt Glowstick renderer internal alpha build

How to Achieve Real Time Path Tracing













Industry Leading Path Tracing Performance without Compromises

- Statistically accurate sampling and filtering
- Texture caching and tiled access
- Progressive rendering
- Unbiased Monte Carlo integration of ray samples
- Path traced global illumination, occlusion, visibility and ray traversal
- Physically accurate reflectance, refraction, transmission, emissivity and caustics
- Energy conserving and preserving illumination models
- Physical area lights and emitters, including HDRI sources
- Camera Depth Of Field for physically accurate focus effects

Full Support for OpenUSD

- Direct interchange with content creation apps
- Fully composable scene hierarchy, cameras, lighting
- Powerful workflows for animated sequences and collaborative editing
- Programmable asset pipeline
- Flexible geometric meshing extensible to SubDivs, NURBS, volumetrics

Hardware Texture Mapping

- Supports OpenImageIO standard
- Cached image buffers
- Tiled and MipMapped textures
- Procedural textures from OpenShadingLanguage
- Direct support for USDImaging
- OpenColorIO for color management
- PTex support



Fully Programmable Shading Pipeline

- USD integrates the ASWF MaterialX standard for material exchange and lookdev
- MaterialX defines USD materials by interconnecting shaders and compute nodes
- Glowstick applies those materials (USDPreviewSurface, USDOpenPBRSurface, USDTexture, etc.) to our USD scene elements
- Glowstick implements BSDFs using OSL
- OSL shaders combine programmable BSDFs, procedural textures, and perform imaging operations
- Glowstick combines OSL shaders into a MaterialX graph
- MaterialX can be parsed to generate OSL shaders
- OSL shaders are compiled to LLVM IR, which is callable from Glowstick



Flexible Rendering API

- Hardware and software emulation on multiple platforms and architectures
- Development tooling on Linux (x86, RISC-V, Arm; Windows coming soon)
- Glowstick can be run as a networked service
- Works with industry standard renderfarm management (Deadline) and job description formats (OpenJD)
 - Supports distributed rendering for rendering animations
 - Supports super-resolution outputs by parallelizing region rendering
- Provides integration to DCC apps through Hydra rendering delegates

Extensible Post Render Pipeline

- Multiple AOV outputs
- Cryptomatte
- Light Path Expressions
- 32b linear output, ACES color managed OpenEXR
- Tone mapping and color grading
- Compositing and post processing workflow
- Compatible with OpenFX and OpenTimelineIO for finishing
- Supports open standards for media storage, broadcast, and archival
- Support USD Asset Management plugins for integration with 3rd party Digital Asset Management systems



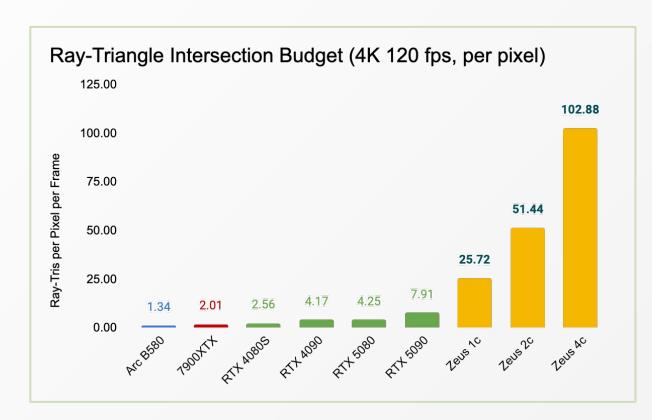


Rendered on an internal alpha build of Glowstick

Zeus Architecture



- For real time path tracing applications, ray per pixel budget improves dramatically
 - Budget includes multiple bounces
 - 4K 120 fps path tracing not achievable on other GPUs
 - Zeus 2c = 8 spp, 5bounces + denoising

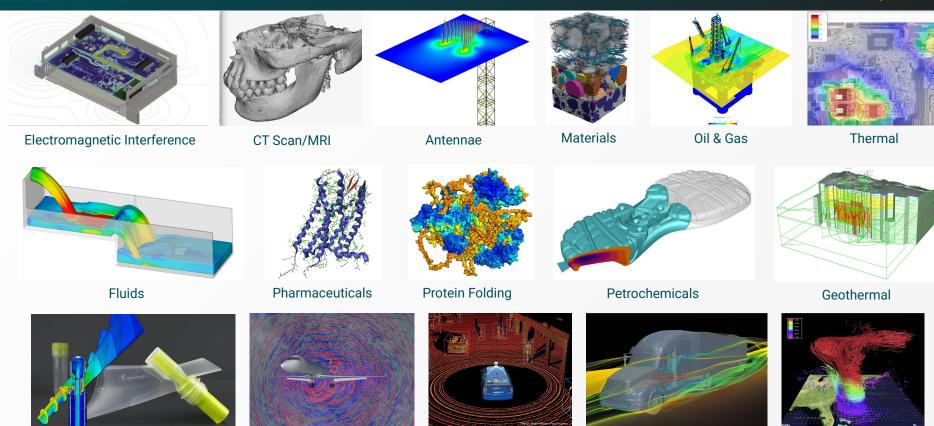


HPC

Simulations are Used Everywhere



Weather



Radar/Lidar

Aerospace

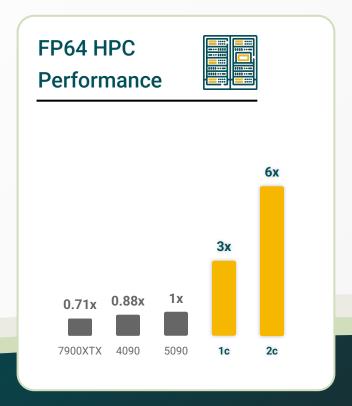
Medical Devices

© 2025 Bolt Graphics Inc. 18

Aerodynamics

HPC Performance Improvements



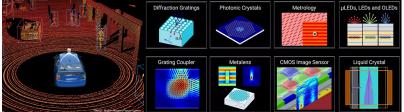


Enhancing Key Workload Performance

O = L TGRAPHICS

Electromagnetic wave (EM) simulation is used to simulate how detectable objects are by radar, silicon photonics to transmit data faster, medical scanners like CT and ultrasound, and EMI/EMC for consumer & industrial electronics.

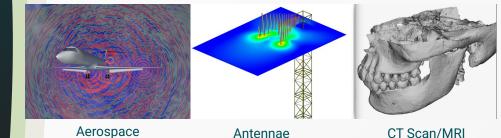




Electromagnetic Interference

Radar/Lidar

Optics



EM Simulation Performance

IEEE-754 FP64 accuracy

300x

0.5x 0.9x 1x
EPYC MI325X B200

© 2025 Bolt Graphics Inc.

20

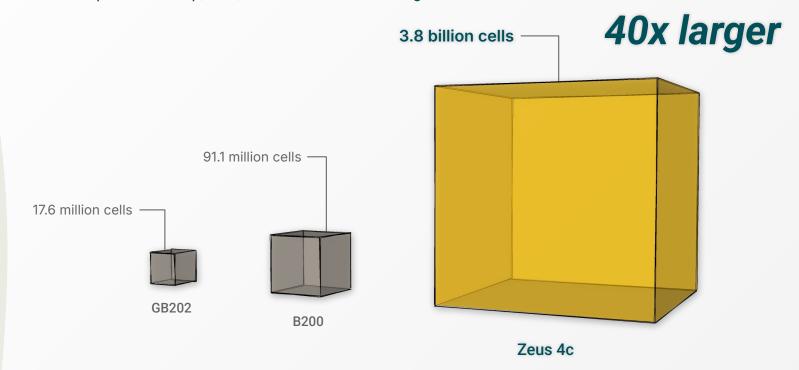
ZEUS

Enhancing Key Workload Performance

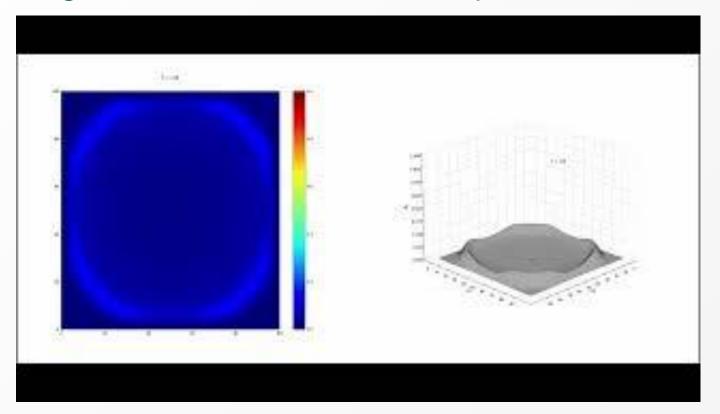


21

Zeus unlocks simulations that were impossible before. Larger simulation spaces enables finer meshes and more complete simulations. Go from simulating a fraction of a photonics chip, lens, or PCB to the entire thing.

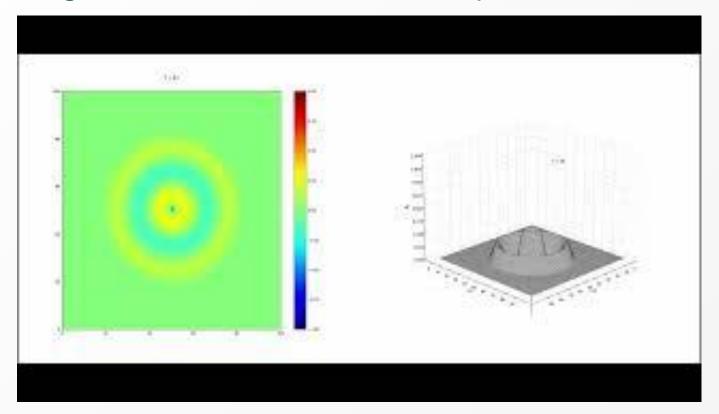






2D Gaussian Pulse Simulation with 3D View simulated on Zeus

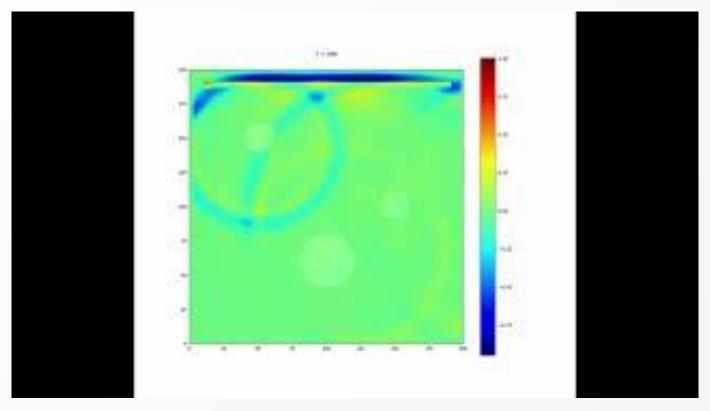




2D Sine Wave Simulation with 3D View simulated on Zeus



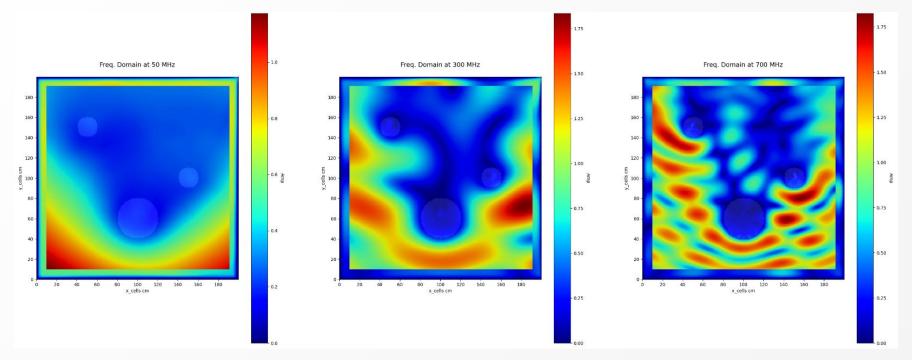
24



Gaussian Plane wave encroaching on 3 dielectric circles (2D simulation on Zeus)

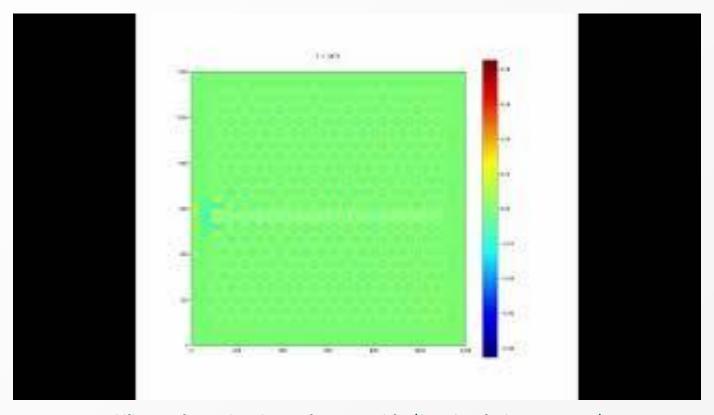
FFT Output





Frequency domain view of the electromagnetic simulation from the previous slide





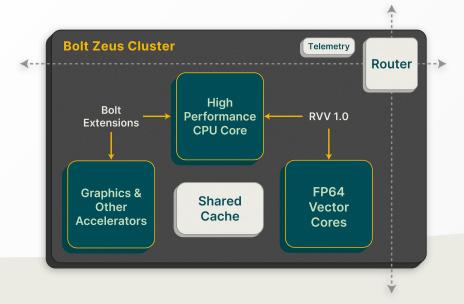
Silicon Photonics Crystal Waveguide (2D simulation on Zeus)



ZEUS ARCHITECTURE



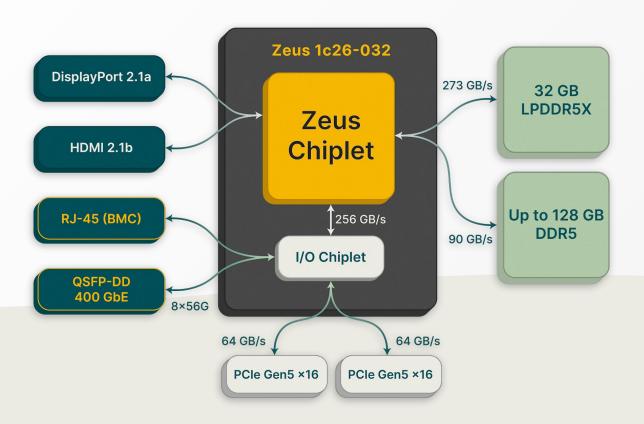
28



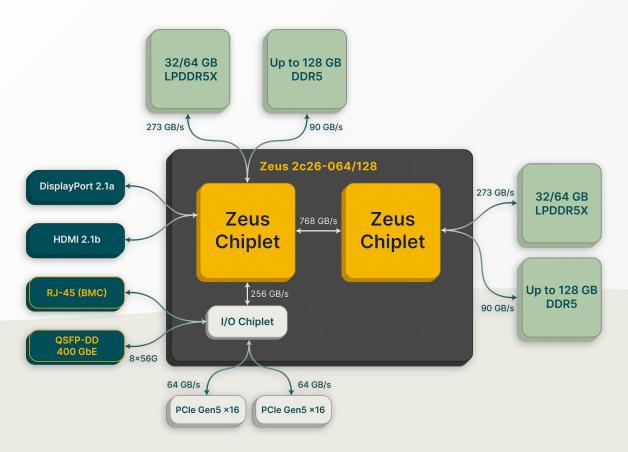
Highly scalable clusters:

- Scalar CPU core (RVA23)
 - High single thread performance
- Vector cores
 - FP64 ALUs
 - RVV 1.0 with slight modifications
- Bolt extensions for other accelerators
- Using RISC-V is smarter than building ISA from scratch - reuse ecosystem momentum and components
- Router for efficient offload of data movement
- Local telemetry

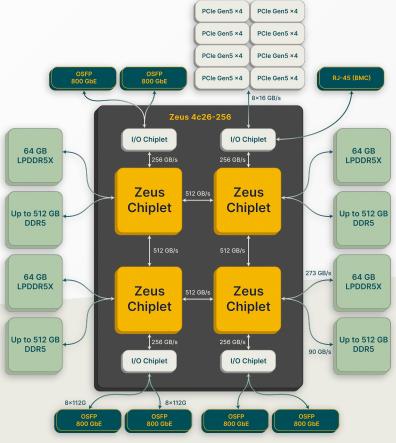








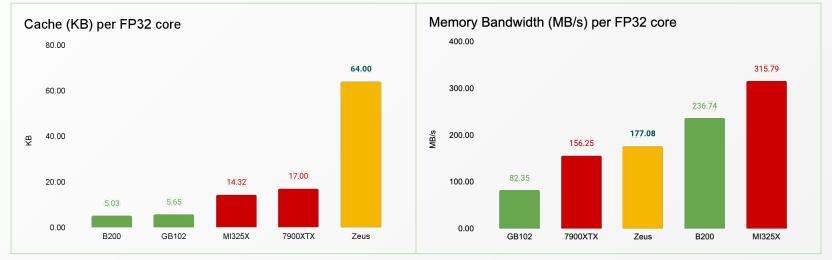




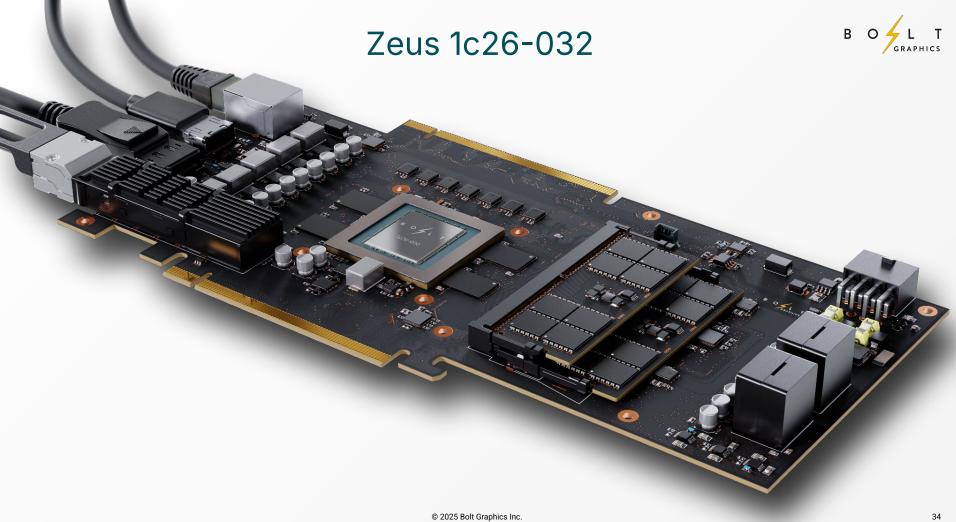
Zeus Architecture

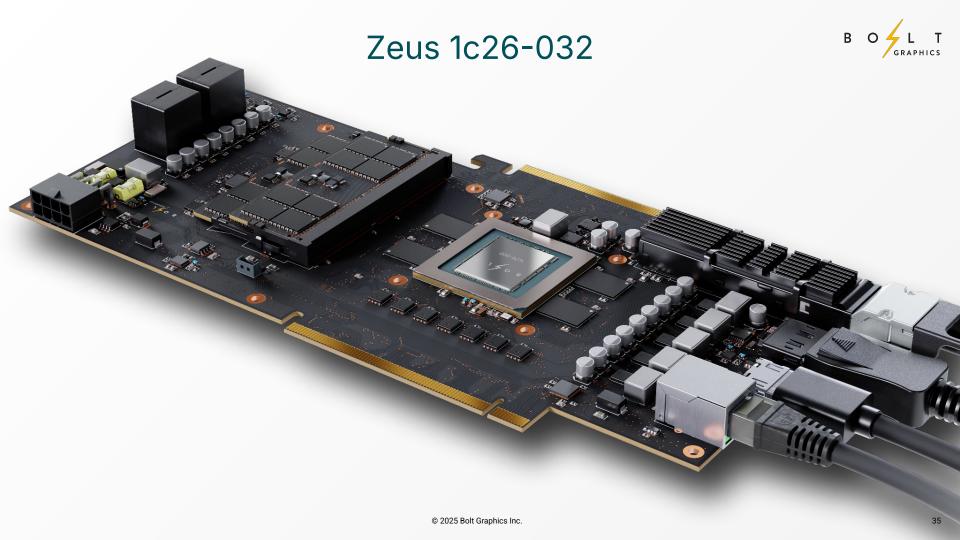


- FP64 native cores (2x FP32, 4x FP16)
 - Memory capacity/bandwidth limited workloads excel on Zeus
 - o Higher ALU occupancy in real workloads not just FP throughput benchmarks
- Extremely high cache to core ratio
 - Keep data closer to compute
- Competitive memory bandwidth to core ratio



ZEUS SKUs





Zeus 1c26-032



	Bolt Zeus 1c26-032	RTX 5080
Board Power	120 W	360 W
FP64 / FP32 / FP16 vector tflops	5 / 10 / 20	0.88 / 56.3 / 56.3
INT16 / INT8 matrix tflops	307.2 / 614.4	439.7 / 879.3
On-chip cache	128 MB	75 MB
Memory	Up to 160 GB @ 363 GB/s 32 GB LPDDR5X 2x DDR5 SO-DIMMs	16 GB @ 960 GB/s
Path Tracing	77 gigarays	17 gigarays
Video Encoding (AV1, H.264/265)	2x 8K60 streams	2x 8K60 streams
I/O	400 GbE (QSFP-DD) 2x PCIe 5.0 x16 DisplayPort 2.1a HDMI 2.1b	PCle 5.0 x16 3x DisplayPort 2.1a HDMI 2.1b

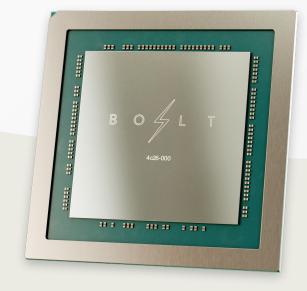
Zeus 2c26-064/128

				/ GRAPHICS
	Bolt Zeus 2c26-064	Bolt Zeus 2c26-128	RTX 4090	RTX 5090
Board Power	250 W		450 W	575 W
FP64 / FP32 / FP16 vector tflops	10 / 20 / 40		1.4 / 90 / 90	1.6 / 105 / 105
FP16 / FP8 matrix tops	614.4 / 1,228.8		714.2 / 1,428.5	818.4 / 1,636.8
On-chip cache	256 MB		88 MB	120 MB
Memory	Up to 320 GB @ 725 GB/s 64 GB LPDDR5X 4x DDR5 SO-DIMMs	Up to 384 GB @ 725 GB/s 128 GB LPDDR5X 4x DDR5 SO-DIMMs	24 GB @ 1 TB/s	32 GB @ 1.8 TB/s
Path Tracing	154 gigarays		17 gigarays	32 gigarays
Video Encoding (AV1, H.264/265)	4x 8K60 AV1 streams		3x 8K60 AV1 streams	
I/O	400 GbE (QSFP-DD) & 2x PCIe 5.0 x16 DisplayPort 2.1a & HDMI 2.1b © 2025 Bolt Graphics Inc.		PCIe 4.0 x16 3x DisplayPort 1.4a HDMI 2.1	PCIe 5.0 x16 3x DisplayPort 2.1a HDMI 2.1b

Zeus 4c26-256



	Bolt Zeus 4c26-256	
Chip Power	500 W	
FP64 / FP32 / FP16 vector tflops	20 / 40 / 80	
FP16 / FP8 matrix tops	1,228.8 / 2,457.6	
On-chip cache	512 MB	
Memory	Up to 2,304 GB @ 1.45 TB/s 256 GB LPDDR5X 8x DDR5 DIMMs	
Path Tracing	307 gigarays	
Video Encoding	8x 8K60 AV1 streams	
I/O	6x 800 GbE (OSFP) 8x PCle 5.0 x4	

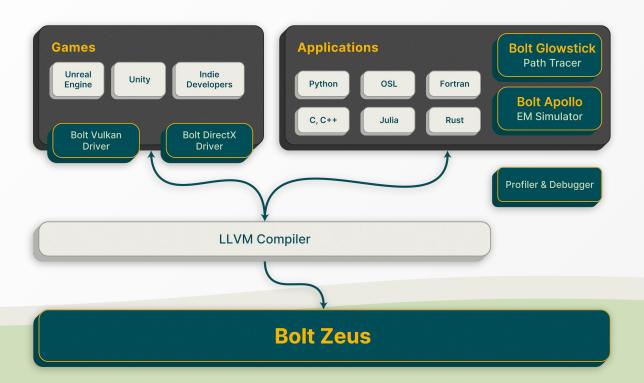




ZEUS SOFTWARE

Zeus Software





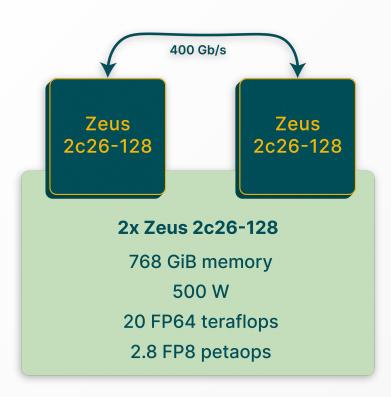
© 2025 Bolt Graphics Inc. 40



ZEUS USAGE SCENARIOS

Zeus Scenario: Small GPU Cluster

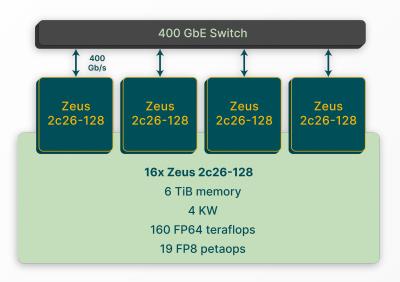




- Directly connect 2 GPUs together over Ethernet or PCIe
- Run the largest physics simulations, Al models, and render the largest scenes with large 768 GiB memory pool

Zeus Scenario: Small Networked GPU Cluster

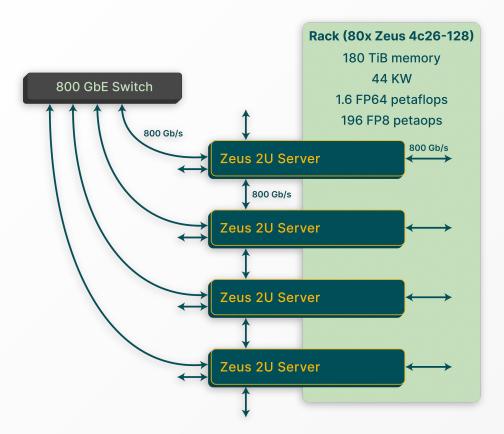




- Connect many GPUs together through Ethernet switch
- Run the largest physics simulations, Al models, and render the largest scenes with huge multi TiBs memory pool

Zeus Scenario: Large Networked GPU Cluster





- Connect many GPUs together through Ethernet switches
 - Also, GPUs are connected to each other in 2D mesh - offloading bandwidth from switch and decreasing neighbor latency
- Run the largest physics simulations, AI models, and render the largest scenes with huge 100s of TiBs memory pool
- 44 KW per rack can be air- or liquid-cooled

© 2025 Bolt Graphics Inc.

Get Connected with Bolt



Meet Us at a Conference for a Live Demo:

- GDC San Francisco, Mar 17-21
- LF Member Summit Napa, March 18-20
- EMA Expo Amherst, April 28-May 2
- GamesBeat Los Angeles, May 19-20
- Computex Taipei, May 19-23
- **ISC** Hamburg, June 10-13

- AWE Long Beach, June 10-12
- Siggraph USA Vancouver, Aug. 10-14
- Hot Chips Palo Alto, Aug. 24-26
- OCP Global Summit San Jose, Oct. 13-16
- Supercomputing St. Louis, Nov. 16-21

Individual User Access

Creators and Researchers interested in Zeus and Glowstick can sign up on our website to be the first in line when Bolt's Early Access Program goes live and to get exclusive merch.

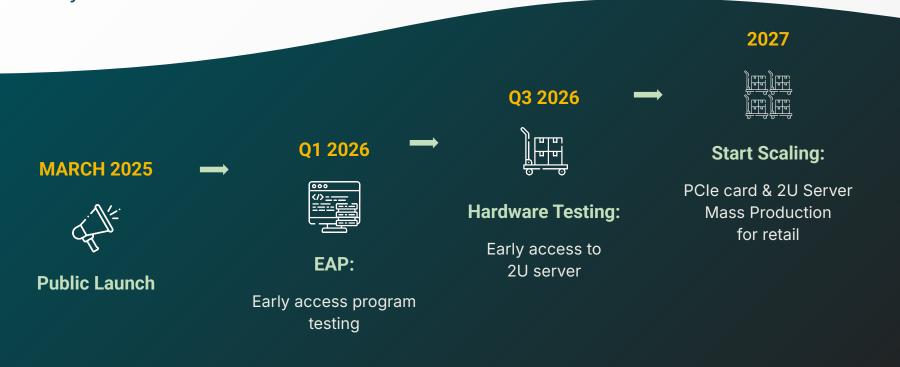
Sign up on our website

© 2025 Bolt Graphics Inc.

Zeus GPU Availability



Partner with Bolt to access developer kits and reserve your Zeus GPUs.





Brad

To truly understand the process and pains of rendering and simulating, we work directly with artists like Matt Curtis to make film quality assets.

This is Brad.v2, and the result of our ongoing collaboration.





