NVIDIA® T1000

PRELIMINARY DOCUMENT – CONTENT SUBJECT TO CHANGE

- NVIDIA Turing Architecture
- 896 NVIDIA® CUDA® Cores
- 4GB GDDR6 Memory
- Up to 160GB/s Memory Bandwidth

- Max. Power Consumption: 50W
- Thermal Solution: Active
- 2.713" H x 6.137" L Single Slot
- Display Connectors: DP 1.4 (4)

NVIDIA GPUs power the world's most advanced desktop workstations, providing the visual computing power required by millions of professionals as part of their daily workflow. All phases of the professional workflow, from creating, editing, and viewing 2D and 3D models and video, to working with multiple applications across several displays, benefit from the power that only a discreet GPU solution can provide. NVIDIA entry level professional GPUs provide advanced features and performance far beyond what is available with integrated GPU solutions – all in a low-profile package, enabling deployment in a wide range of small form factor workstations. Step up to the power of an NVIDIA discreet professional GPU.

As part of the NVIDIA RTX family of professional GPUs, the NVIDIA T1000 provides the performance, features, reliability, and support that enterprise customers expect from NVIDIA professional visual computing solutions.

Performance Features

Turing GPU Architecture
Based on state-of-the-art 12nm FFN (FinFET NVIDIA) high-performance manufacturing process customized for NVIDIA to incorporate 896 CUDA cores, the NVIDIA T1000 GPU is the most powerful Single Slot professional solution for CAD, DCC, financial service industry (FSI) and visualization professionals in general looking to reach excellence performance in a compact and efficient form factor. The Turing GPU architecture enables the biggest leap in computer real-time graphics rendering since NVIDIA's invention of programmable shaders in 2001.

Advanced Shading Technologies
The Turing GPU architecture features the following new advanced shader technologies.

Mesh Shading: Compute-based geometry pipeline to speed geometry processing and culling on geometrically complex models and scenes. Mesh shading provides up to 2x performance improvement on geometry-bound workloads.

Variable Rate Shading (VRS): Gain rendering efficiency by varying the shading rate based on scene content, direction of gaze, and motion. Variable rate shading provides similar image quality with 50% reduction in shaded pixels.

Texture Space Shading: Object/texture space shading to improve the performance of pixel shader-heavy workloads such as depth-of-field and motion blur. Texture space shading provides greater throughput with increased fidelity by reusing pre-shaded texels for pixel-shader heavy VR workloads.

Advanced Streaming Multiprocessor (SM) Architecture

Combined shared memory and L1 cache improve performance significantly, while simplifying programing and reducing the tuning required to attain best application performance. Each SM contains 96 KB of L1 shared memory, which can be configured for various capabilities depending on compute or graphics workload. For compute cases, up to 64KB can be allocated to the L1 cache or shared memory, while graphics workload can allocate up to 48 KB for shared memory; 32 KB for L1 and 16KB for texture units. Combining the L1 data cache with the shared memory reduces latency and provide higher bandwidth.

Higher Speed GDDR6 Memory

Built with Turing's vastly optimized 4GB GDDR6 memory subsystem for the industry's fastest graphics the NVIDIA T1000 features 4GB of frame buffer capacity and 160 GB/s of peak bandwidth for double the throughput from previous generation. NVIDIA T1000 is the ideal platform for 3D professionals and high demanding with vast arrays of datasets and multi display environments.

Single Instruction, Multiple Thread (SIMT)

New independent thread scheduling capability enables finer-grain synchronization and cooperation between parallel threads by sharing resources among small jobs.

Mixed-Precision Computing

Double the throughput and reduce storage requirements with 16-bit floating point precision computing to enable the training and deployment of larger neural networks. With independent parallel integer and floating-point data paths, the Turing SM is also much more efficient on workloads with a mix of computation and addressing calculations.
software-based x264 encoders.

Graphics Preemption

Pixel-level preemption provides more granular control to better support time-sensitive tasks such as VR motion tracking.

Compute Preemption

Preemption at the instruction-level provides finer grain control over compute tasks to prevent long-running applications from either monopolizing system resources or timing out.

H.264 and HEVC Encode/Decode Engines

Deliver faster than real-time performance for transcoding, video editing, and other encoding applications with two dedicated H.264 and HEVC encode engines and a dedicated decode engine that are independent of 3D/compute pipeline.

NVIDIA GPU BOOST 4.0

Automatically maximize application performance without exceeding the power and thermal envelope of the card. Allows applications to stay within the boost clock state longer under higher temperature threshold before dropping to a secondary temperature setting base clock.

Image Quality

Full-Scene Antialiasing (FSAA)

Dramatically reduce visual aliasing artifacts or "jaggies" with up to 64X FSAA (128X with SLI) for unparalleled image quality and highly realistic scenes.

32K Texture and Render Processing

Texture from and render to 32K x 32K surfaces to support applications that demand the highest resolution and quality image processing.

Display Features

NVIDIA® NVIDIA® Mosaic Technology

Transparently scale the desktop and applications across up to 4 GPUs and 16 displays from a single workstation while delivering full performance and image quality.

DisplayPort 1.4a

Support up to four 5K monitors @ 60Hz, or dual 8K displays per card. NVIDIA T1000 supports 4x displays natively, HDR color for 4K @ 120Hz for 10/12b HEVC decode and up to 4K @ 60Hz for 10b HEVC encode. Each DisplayPort connector is capable of driving ultra-high resolutions of 4096x2160 @ 120 Hz with 30-bit color.

NVIDIA® RTX™ Desktop Manager

NVIDIA RTX™ Desktop Manager software allows you to manage single or multi-monitor workspaces with ease, giving you maximum flexibility and control over your display real estate and desktops.

OpenGL Quad Buffered Stereo Support

Provide a smooth and immersive 3D Stereo experience for professional applications.

NVIDIA® Enterprise Management Tools

Maximize system uptime, seamlessly manage wide-scale deployments and remotely control graphics and display settings for efficient operations.

Software Support

NVIDIA® CUDA® Parallel Computing Platform

Natively execute standard programming languages like C/C++ and Fortran, and APIs such as OpenCL, OpenACC and Direct Compute to accelerates techniques such as ray tracing, video and image processing, and computation fluid dynamics.

Unified Memory

Featuring concurrent execution of floating point and integer operations, adaptive shading technology, and a new unified memory architecture with twice the cache of its predecessor, Turing shaders enable awesome performance increases for today's professional applications.

Specifications

GPU Architecture	NVIDIA Turing Architecture
CUDA Parallel Processing cores	896
Single-Precision Performance	Up to 2.5 TFLOPS
Frame Buffer Memory	4 GB GDDR6
Memory Interface	128-bit
Memory Bandwidth	Up to 160 GB/s
Max Power Consumption	50 W
Graphics Bus	PCI Express 3.0 x16
Display Connectors	DP 1.4 (4)
Form Factor	2.713 inches H x 6.137 inches L Single Slot
Product Weight	132.6g
Thermal Solution	Active