NVIDIA RTX 6000 Ada Generation Desktop Workstation Messaging Document

Initial version: 7/19/2022
Document Owner: Stacy Ozorio

Revision History

(All changes must be approved by the Document Owner)

Version	Revision Date	Edited By
0.1	7/19/2022	Stacy Ozorio
0.2	7/26/2022	Stacy Ozorio
0.3	7/27/2022	Stacy Ozorio
0.4	7/29/2022	Stacy Ozorio
0.5	8/1/2022	Stacy Ozorio
0.6	8/2/2022	Stacy Ozorio
0.7	8/5/2022	Stacy Ozorio
0.8	8/15/2022	Stacy Ozorio, Allen Bourgoyne, Andrew Rink
0.9	8/18/2022	Stacy Ozorio, Sepi Motamedi
1.0	8/19/2022	Stacy Ozorio
1.1	8/25/2022	Stacy Ozorio, Himanshu Iyer
1.2	9/8/2022	Stacy Ozorio, Rick Champagne
1.3	9/20/2022	Stacy Ozorio
1.4	9/28/2022	Stacy Ozorio
1.5	10/7/2022	Sepi Motamedi
1.6	11/9/2022	Stacy Ozorio
1.7	11/16/2022	Stacy Ozorio, James Kim

Performance numbers in this document may be subject to change until product availability.

Table of Contents

NVIDIA RTX 6000 Ada Generation Desktop Workstation Messaging

NVIDIA RTX Key Message

Tagline: Performance for Endless Possibilities

In 2018, the NVIDIA RTX™ platform revolutionized professional visual computing with the NVIDIA Turing™ GPU architecture. It combined the power of advanced graphics, compute, and AI to bring real-time ray tracing to reality years before previously thought possible. NVIDIA RTX technology enabled artists, designers, scientists, and researchers to work with photorealistic images in real time, fundamentally changing the way they work.

Since the introduction of NVIDIA RTX technology, the world and how people work has changed drastically. Professional workflows need to enable a distributed and remote workforce with the ability to do any work from anywhere. Artists are creating higher-resolution, hyper realistic content with ever-increasing demand for differentiated and visually compelling content. Designers and engineers are striving to create more complex and efficient designs to meet customer demands in highly supply-constrained environments. Scientists, researchers, and medical professionals are faced with incredible challenges that require the rapid development of solutions on a global scale.

With the introduction of the NVIDIA RTX 6000 Ada Generation desktop GPU, based on the NVIDIA Ada Lovelace architecture, we are providing incredible new levels of performance for professionals to take on the challenges of today and be ready for the evolving, intensive demands of tomorrow.

Copy Blocks for RTX 6000

100 Words

The way people work is undergoing a drastic change with distributed teams and remote workers as the new normal. Artists are facing an ever-increasing demand for differentiated and visually compelling content. Designers and engineers are striving to create more complex and efficient designs in highly supply-constrained environments. Scientists, researchers, and medical professionals are faced with incredible challenges that require the rapid development of solutions on a global scale.

The NVIDIA RTX™ 6000 Ada Generation is designed to meet the challenges of today's professional workflows. Built on the NVIDIA Ada Lovelace architecture, the RTX 6000 combines 142 third-generation RT Cores, 568 fourth-generation Tensor Cores, and 18,176 CUDA® cores with 48GB of graphics memory to deliver the next generation of AI graphics and petaflop inferencing performance for unprecedented speed-up of rendering, AI, graphics, and compute workloads. RTX 6000-powered workstations provide what you need to succeed in today's ultra-challenging business environment.

50 Words

The NVIDIA RTX™ 6000 Ada Generation delivers the features, capabilities, and performance to meet the challenges of today's professional workflows. Built on the NVIDIA Ada Lovelace architecture, the RTX 6000 combines 142 third-generation RT Cores, 568 fourth-generation Tensor Cores, and 18,176 CUDA® cores with 48GB of graphics memory for unprecedented rendering, AI, graphics, and compute workload performance. RTX 6000-powered workstations provide what you need to succeed in today's ultra-challenging business environment.

25 Words

The NVIDIA RTX™ 6000 Ada Generation combines the latest-generation RT Cores, Tensor Cores, and CUDA® cores with 48GB of graphics memory to deliver the ultimate desktop performance designers, engineers, and artists need to propel innovation forward.

Support/Benefit Messages

Value for End-Users

- **Incredible Application Performance**
  - Experience fast, interactive performance—powered by the latest NVIDIA Ada Lovelace architecture-based GPU—with ultra-fast, onboard graphics memory technology and optimized software drivers for professional applications. The Shader Execution Reordering (SER) system allows on-the-fly organization and reordering workloads by grouping similar performing threads for the SM and RT Core to operate more efficiently.

- The NVIDIA RTX 6000 Ada Generation includes 142 RT Cores to accelerate photorealistic ray-traced rendering with up to 2X faster ray-triangle intersection throughput than the previous generation. The Opacity Micromap (OMM) engine unit enables up to 2X faster alpha traversal ray tracing for transparent surfaces like foliage, particles, and fences. The Displaced Micro-Mesh (DMM) engine unit optimizes the Bounding Volume Hierarchy (BVH) structure by reducing BVH build time by 10X and reducing BVH storage by 20X compared to traditional rendering of complex objects with high levels of geometric detail.
- With 568 Tensor Cores to accelerate AI workflows, the RTX 6000 provides the performance necessary for AI development and training workloads. Incredible inferencing performance and enterprise-class stability and reliability make RTX 6000-powered desktop workstations ideal for professional AI training and inferencing deployments.
- Ensure hardware compatibility and stability through NVIDIA support of the latest OpenGL, DirectX, Vulkan, and CUDA® standards, deep independent software vendor (ISV) developer engagements, and certification with over 100 professional software applications.
- NGC™ support gives engineers, researchers, and data scientists access to containers for the top deep learning frameworks that are tuned, tested, certified, and maintained by NVIDIA, as well as third-party managed high-performance computing (HPC) containers, NVIDIA HPC visualization containers, and partner applications.

- **Rich, Expansive Visual Workspace**
  - Experience stunning imagery through movie-quality, anti-aliasing techniques, high-dynamic range (HDR) color support, higher refresh rates, and up to 8K screen resolution at 60Hz from a single cable with the DisplayPort 1.4a standard.
  - Enhance your desktop workspace experience with NVIDIA RTX™ Desktop Manager[1] and NVIDIA Mosaic technology. Work across four displays on every NVIDIA RTX professional card with the intuitive placement of windows, multiple virtual desktops, and user profiles.
  - Use advanced multi-display technologies like Quadro® Sync II, NVIDIA Mosaic, and Warp and Blend to synchronize images and scale resolution on a display surface with multiple projectors or screens.

Value for IT Administrators

- **Proven Reliability**

[1] Product formerly known as NVIDIA Quadro View has undergone a brand transition.

- Experience higher-quality products driven by power-efficient hardware and components selected for optimum operational performance, durability, and longevity.
- Enjoy maximum uptime thanks to exhaustive product testing with leading OEMs and system integrators that simulates the most demanding real-world conditions.

- **Easy Manageability**
  - Remotely monitor and manage NVIDIA professional products in your enterprise by integrating the NVIDIA Enterprise Management Toolkit (NVWMI) in your IT asset management framework.
  - Scale up NVIDIA RTX Enterprise driver[2] deployment to hundreds of workstations using NVWMI's powerful driver installer.
  - Simplify software driver deployment through a regular cadence of long-life, stable driver releases based on a robust feature-development and quality-assurance process.

Value Proposition

Component	Copy
FOR (Target Audience)	Creative and technical professionals, software developers, and data scientists across manufacturing, media and entertainment (M&E), medical, architecture, engineering, construction, and operations (AECO), oil and gas, scientific visualization, and AI/deep learning market segments.
WHO (Statement of Need or Opportunity)	Professionals who need the most powerful real-time rendering, graphics, AI, VR, and compute solution available in a desktop workstation platform.
(Product Name) IS A (Product Category)	The NVIDIA RTX 6000 combines the highest levels of rendering, visualization, AI, and compute performance available in a professional graphics card.
THAT (Compelling Reason to Buy)	Provides the capability for massive graphics and visualization, real-time photorealistic rendering, immersive VR environments, powerful compute, and creation and deployment of AI solutions.

[2] Prior Quadro Optimal Driver for Enterprise (ODE) driver has undergone a brand transition, effective 12/14/20.

UNLIKE (Primary Alternatives)	AMD Radeon Pro graphics
ONLY (Primary Differentiator, Reason to Believe)	The NVIDIA RTX 6000 delivers the most powerful real-time rendering, graphics, VR, compute, and deep learning solutions available in a professional desktop workstation.

Additional Messaging

Products	Support Message	Proof Points
NVIDIA RTX 6000	The new NVIDIA RTX 6000	• CUDA cores: The NVIDIA Ada Lovelace architecture-based CUDA cores bring more than 2X the single-precision floating-point (FP32) throughput compared to the previous generation, providing significant performance improvements for graphics workflows such as 3D model development and compute for workloads such as desktop simulation for computer-aided engineering (CAE). • RT Cores: Third-generation RT Cores provide up to 2X the throughput of the previous generation, and the ability to concurrently run ray tracing with either shading or denoising capabilities. This accelerates renders for M&E content creation, AECO design evaluations, and manufacturing virtual prototyping. Third-generation RT Cores deliver up to 2X the ray tracing performance over the previous generation, delivering groundbreaking performance for photorealistic rendering. Enhanced RT Cores combined with new Shader Execution Reordering (SER) technology dynamically reorder inefficient workloads, dramatically improving shader performance to accelerate end-to-end ray-traced image rendering performance.

Performance numbers in this document may be subject to change until product availability.

		• Tensor Cores: Fourth-generation Tensor Cores provide up to 2X faster AI training performance than the previous generation with FP16 precision. Support for the new FP8 data format for inferencing provides more than 4X faster performance than the previous generation and reduces data memory usage by half (compared to FP16 data format) • Encode/decode engines: The RTX 6000 includes three video encode engines and three decode engines, including support for the AV1 video format and the performance required for multistream video applications for security and video serving. • PCIe Gen 4: The RTX 6000 supports PCI Express Gen 4 (PCIe Gen 4), which doubles the bandwidth of PCIe Gen 3 from 15.75 gigabytes per second (GB/s) to 31.5GB/s for x16 connection, improving data transfer speeds from CPU memory for data-intensive tasks such as AI, data science, and creating 3D models from large datasets. Faster PCIe performance also accelerates GPU direct memory access (DMA) transfers, providing faster video data transfers from GPUDirect® for Video-enabled devices and faster input/output (I/O) with GPUDirect Storage. • 48GB of GPU memory: The RTX 6000 features 48GB of GDDR6 memory, providing the memory needed for rendering, data science, engineering simulation, and other GPU memory-intensive applications. With greater memory bandwidth than the previous generation, the RTX 6000 can move data between the GPU and GPU memory faster, resulting in better graphics, compute, and rendering performance. • Virtualization-ready[3]: Support for NVIDIA virtual GPU (vGPU) software allows a personal workstation to be repurposed into multiple high-performance virtual workstation instances. This

[3] Virtualization support for the RTX 6000 GPU will be available in an upcoming NVIDIA virtual GPU (vGPU) release.

		enables remote users to share resources to drive high-end design, AI, and compute workloads. • Multi-display: Four DisplayPort 1.4a connectors, support for multiple 8K monitors, Quadro Sync, NVIDIA Mosaic, and Warp and Blend enable massive immersive environments for cave automatic virtual environments (CAVE), video walls, and location-based entertainment deployments. • AR/VR: Support for the latest high-resolution head-mounted display (HMD) devices, high-performance graphics, and large 48GB of GPU memory enables the creation of incredibly immersive augmented reality (AR) and virtual reality (VR) experiences for training, product validations, building walkthroughs, and compelling entertainment. • Form factor and power efficiency: With a dual-slot, power-efficient design, the RTX 6000 fits into a wide range of workstation chassis, providing professionals with a generous selection of compatible workstations from worldwide OEM vendors. • Motion BVH (bounding volume hierarchy): Hardware-accelerated rendering of motion blur—a common cinematic effect that is difficult to render—means artists no longer need to rely on traditional methods of using motion vectors to achieve motion blur. Motion vectors give the artist flexibility to adjust motion blur in post but require visual fixes for reflections and translucency. • NVIDIA Deep Learning Super Sampling (DLSS) 3.0: The Ada Lovelace GPU architecture features a new Optical Flow Accelerator and AI-based DLSS Super Resolution with DL denoiser that boost DLSS 3.0's frame rates up to 4X compared to the previous version while maintaining or exceeding native image quality.

		• NVIDIA RTX Broadcast Engine: The NVIDIA RTX Broadcast Engine transforms offices into broadcast studios, upgrading standard webcams and microphones into premium, smart devices with the power of AI. Improve the video and audio quality of livestreams with AI capabilities such as virtual backgrounds, webcam auto frame, and microphone noise removal. With dedicated AI processors called Tensor Cores on NVIDIA RTX GPUs, the AI networks can run in real time alongside applications.
	The NVIDIA RTX 6000 display technology	• Display support: A single RTX 6000 can support up to four 5K (5120 x 2880) at 60Hz displays. Display support also includes the ability to drive two 8K displays. The RTX 6000 supports DisplayPort 1.4a, which enables a single cable to drive an 8K display at 60Hz using Display Stream Compression (DSC). • Advanced multi-display technologies such as Quadro Sync, NVIDIA Mosaic, and Warp and Blend provide a robust way to perform image synchronization and resolution scaling of a synchronized display surface with multiple projectors or displays. • NVIDIA RTX Desktop Manager[1] offers intuitive window management across multiple desktops and applications that let users spread work across a single large display or multiple displays to optimize workflows. • The RTX 6000 provides the display performance required by the latest generation of higher-resolution HMDs.

		• Enhance the analysis and visualization of design, engineering, and digital-content creation workflows with stereoscopic 3D display from a wide range of stereoscopic display devices.
	Reliability is core to all NVIDIA products. NVIDIA is fundamentally designed to deliver the peace of mind that customers expect as they focus on doing what they do best — changing the way we look at the world.	• NVIDIA workstation products are designed, built, tested, and supported exclusively by NVIDIA—not by third-party add-in board vendors—to ensure maximum uptime. Components rated for higher performance, quality, and longevity are designed to deliver the highest level of product reliability. • Long-life, unified software drivers for the enterprise undergo extensive testing. This includes subjecting a broad range of workstations and business PCs carrying various NVIDIA workstation products to tens of thousands of machine hours of stress tests and thousands of hours of attended platform testing. • Rigorous joint-qualification programs with leading OEMs like HP, Dell, and Lenovo require that a comprehensive set of NVIDIA workstation system configurations pass an exhaustive battery of hardware and software validation tests. This ensures that the final product meets customers' expected reliability and quality standards. • NVIDIA workstation products have earned application certifications from more than 100 independent software vendors (ISVs). Comprehensive compatibility testing and performance tuning with leading third-party ISVs across a broad range of industries ensure optimal and stable performance with current versions of their applications and support for the latest features.

	NVIDIA IT manageability	• The NVIDIA Enterprise Management Toolkit (NVWMI) offers IT administrators a robust set of asset management capabilities, including end-user display configuration and GPU monitoring. • The toolkit's NVIDIA driver installer provides a flexible, convenient way to scale up the deployment of driver updates to hundreds of workstations from anywhere on the company network.
	NVIDIA workstation solutions provide unique software tools and utilities designed to enhance professional workflows.	• NVIDIA RTX Experience™[4] provides a full suite of productivity tools on NVIDIA desktop workstations. Features like automatic alerts for the latest NVIDIA RTX Enterprise Driver[1] ensure your system is always optimized for performance. With native desktop recording in up to 8K and tools like Instant Replay, professionals can easily capture workflows, tutorials, and trainings and upload to third-party sites like Twitch, broadcast live, or stream to NVIDIA SHIELD for easy collaboration. NVIDIA RTX Experience also provides a great gaming experience with access to features like optimal playable settings (OPS) and NVIDIA Ansel and Freestyle. • NVIDIA RTX Desktop Manager[2] helps professionals organize single- or multi-monitor workspaces to boost productivity. Features like flexible window snapping, user profiles, windows management, and hotkey support deliver maximum control over desktops and display real estate.

[4] Product formerly known as NVIDIA Quadro Experience, rebrand effective 4/21/2021.

Industry Messaging Points (Key Industries)

Manufacturing / Product Development

- The NVIDIA RTX™ 6000 Ada Generation features an industry-leading 48GB of GDDR6 memory, which enables product designers, engineers, and engineering analysts in the Manufacturing industry to easily work with large datasets and tackle complex workflows across all product design and engineering applications.
- Third-generation RT Cores in the NVIDIA RTX 6000 provide up to 2X faster ray-triangle intersection throughput of the previous generation and enable concurrent ray tracing and shading. Design reviews can be accelerated to review high-fidelity visualizations with physically accurate materials, lighting, and reflections in real time. This improved performance unlocks the ability to evaluate new product concepts rapidly and create stunning marketing content directly from computer-aided design (CAD) geometry.
- Extend the power of engineering simulation throughout the design process with accelerated real-time structural, thermal, and fluid flow analysis software tools. Engineers can instantly evaluate the impact of design changes to iterate at the speed of light and revolutionize the performance of their designs. The NVIDIA Ada Lovelace GPU architecture-based CUDA cores bring more than 2X the single-precision floating point throughput and greater than 2X improved power efficiency compared to previous generations, unlocking new potential for the industry.
- Support for the latest high-resolution HMD devices, high-performance graphics, and huge 48GB of GPU memory enable the creation of incredibly immersive virtual reality experiences to accelerate product development by improving design reviews, training, validation, and production planning.

Architecture, Engineering, Construction, and Operations (AECO)

- The NVIDIA RTX™ 6000 Ada Generation features an industry-leading 48GB of GDDR6 memory, which enables AECO teams to work with massive building information modeling datasets. Architects, designers, and engineers can take advantage of the latest advanced visualization technologies such as ray-traced rendering, virtual reality, and reality capture, that are transforming how the AECO industry works.
- The NVIDIA RTX 6000 Ada Generation includes 142 RT Cores to accelerate photorealistic ray-traced rendering with up to 2X faster ray-triangle intersection throughput than the previous generation. Design teams and architectural visualization specialists can visualize high-fidelity renders in real time using physically based materials and lighting, enabling them to rapidly iterate on ideas to deliver the most accurate and compelling images.
- Support for the latest high-resolution HMD devices, high-performance graphics, and huge 48GB of GPU memory enable the creation of incredibly immersive virtual reality experiences for better design, design reviews, training, and virtual construction rehearsal workflows.

Media and Entertainment (M&E)

- The NVIDIA RTX™ 6000 is NVIDIA's most advanced professional GPU for Media & Entertainment applications. Featuring an industry-leading 48GB of GDDR6 memory, the NVIDIA RTX 6000 enables studios to work in real time with complex geometry and high-resolution textures. Artists can work interactively with massive photorealistic environments made up of hundreds of billions of polygons.
- Iterate more, render faster, and gain distinct advantages over CPU rendering, from previsualization through final frames. Third-generation RT Cores in the NVIDIA RTX 6000 Ada Generation improve ray-tracing performance by up to 2X. Photorealistic ray tracing can now be achieved in real time, enabling higher-fidelity workflows, from interactive rendering to real-time virtual production.
- Dramatically accelerate production pipelines with fourth-generation Tensor Cores, which provide more than 4X the training with FP8 precision throughput over the previous generation with zero code changes. Many creative applications have adopted AI-accelerated features that can aid artists with time-consuming tasks such as upscaling or retiming video, denoising renders, and enhancing audio. Apps like Autodesk Flame, Adobe Premiere Pro®, and Blackmagic Design DaVinci Resolve already take advantage of Tensor Cores to accelerate content creation tasks.
- Broadcasters benefit from NVIDIA RTX 6000's third-generation RT Cores. Improved ray-tracing performance by up to 2X and 48 GB of GPU memory enables real-time composition of augmented reality, ultra-high-resolution video, and stunning graphics on next-generation virtual studios.
- 8th Generation Video Encoders, now with three Encoders and AV1, enabling you to export video with more than 3X throughput vs prior gen, and broadcast video with more quality thanks to AV1 providing up to 40% more efficient compression over H.264.
- Take video processing and transcoding to the next level with NVIDIA RTX 6000's three video encode and decode engines, including support for the AV1 video format and the performance required for multi-stream video applications.
- Fourth-generation Tensor Cores provide up to 2X faster AI training performance with FP16 precision than the previous generation, accelerating AI workloads in broadcast including content discovery, real-time subtitling and translation, and content and viewer analytics.

Energy

- With 91.1 TFLOPS of FP32 compute performance and 1457.0 TFLOPS of FP8 Tensor performance with structural sparsity, the NVIDIA RTX 6000 enables geoscientists to speed up their visualization workflows and gain faster insight into subsurface data.
- Fourth-generation Tensor Core technology with FP8 precision for deep learning enables the RTX 6000 to drive the latest AI-augmented exploration and production software workflows faster than ever.

Performance numbers in this document may be subject to change until product availability.

- Powerful performance enables the RTX 6000 to drive the most immersive VR environments using the latest generation of high-resolution HMDs for realistic immersive training, well planning, and exploration.

Scientific Visualization

- With real-time ray-tracing capabilities, the NVIDIA RTX™ 6000 Ada Generation enables the creation of the most vivid GPU-accelerated renders, providing valuable insights to scientists and researchers faster than ever.
- Dedicated deep learning hardware in the RTX 6000 helps researchers create new AI-augmented technologies for solving the most difficult problems or deploy AI-integrated solutions to drive discoveries and push the boundaries of exploration.
- VR enables scientists to gain insights only available in immersive environments. The RTX 6000 supports the resolutions required by next-generation HMDs.
- The NVIDIA Ada Lovelace architecture provides the RTX 6000 with the power to visualize massive datasets at resolutions of up to 8K.

AI/Data Science

- Fourth-generation Tensor Cores support the FP8 precision that provides more than 4X the training performance of the previous generation for faster model training.
- PCIe Gen 4 provides 2X the bandwidth of Gen 3, speeding up data transfers from CPU and storage, accelerating data prep, and training dataset loading.
- The NVIDIA RTX™ 6000 Ada Generation provides 48GB of GPU memory to work with large datasets for data prep and model training.

Performance numbers in this document may be subject to change until product availability.