HOW
NVIDIA QUADRO RTX
IS TRANSFORMING ARCHITECTURAL VISUALIZATION

Today's challenging architectural projects demand powerful graphics cards. Whether you're rendering high-resolution marketing images in V-Ray, generating real-time walkthroughs in Enscape, or collaborating on room-scale VR experiences in Unreal Engine, your GPU lies at the heart of your production pipeline. With the right workstation graphics card, architectural professionals can explore larger CAD and BIM data sets, iterate and validate designs more quickly, and deliver more content to their clients in less time.

Available worldwide through PNY Technologies, NVIDIA® Quadro® RTX™ GPUs provide the power you need to tackle such complex tasks. For years, AEC professionals have turned to NVIDIA hardware for its combination of compute performance, graphics memory, and support for industry-standard software. Now, with their revolutionary Turing architecture and NVIDIA's new dedicated RT ray tracing Cores, the NVIDIA Quadro RTX series of cards – including the NVIDIA Quadro RTX 4000, 5000, 6000 and 8000 – are opening up new workflows in architectural design and visualization.

KEY NVIDIA QUADRO RTX BENEFITS: SHARED GPU MEMORY, AI ACCELERATION AND REAL-TIME RAY TRACING

The Turing architecture has several features that benefit graphically intensive tasks. One is the second generation of NVLink, NVIDIA's high-speed GPU interconnection technology. With 48GB of error-correcting GDDR6 RAM, the NVIDIA Quadro RTX 8000 has more on-board graphics memory than any other commercially available workstation card. NVLink further enables users to split scene data across two GPUs, creating a potential pool of 96GB of shared memory: enough to fit even the largest architectural scene into the frame buffer during GPU rendering.

Turing cards also have three types of processor Cores tailored to specific GPU computation tasks. CUDA Cores, introduced in previous generations of NVIDIA hardware, accelerate general compute operations, and are now widely supported in GPU rendering applications, including many of the key third-party renderers for 3ds Max.

In addition, the Turing architecture's Tensor Cores are designed specifically for the Tensor/matrix operations used in deep learning. In architectural pipelines, they accelerate AI-driven workflows like generative design and denoising: NVIDIA AI is integrated into key rendering applications, including V-Ray and Corona Renderer, enabling users to generate clean, noise-free images as they edit more quickly and efficiently than before.

The third type of processor core is transformative. The NVIDIA Quadro RTX cards' new RT Cores accelerate ray tracing, enabling a single GPU to render images with physically accurate shadows, reflections and refractions in real time. On the most challenging jobs, they provide an extra performance boost, right when you need it most.

Production renders go 40% faster on average on NVIDIA RTX GPUs using V-Ray GPU's new RTX engine. Scene courtesy of Bertrand Benoit.

POWERING A NEW GENERATION OF RAY TRACING APPLICATIONS

"There are two things that determine the speed of a GPU for rendering: number of CUDA Cores, and clock speed. The RT Cores are now a kicker on top of that,"

Phil Miller,
VP of Product Management at Chaos Group.

One of the first companies to actively take advantage of the Turing architecture in its applications, Chaos Group uses NVIDIA's OptiX API to target the cards' RT Cores in V-Ray GPU's new RTX render engine. RTX acceleration is already available in V-Ray Next for 3ds Max and Maya, and the company plans to roll out active RTX support in other V-Ray applications as well.

"V-Ray GPU has been rendering on NVIDIA RTX cards since they first shipped, but if you add in the RT Cores, that's when you get the big performance boost: an average of 40% speed increase,"

Phil Miller,
VP of Product Management at Chaos Group.

In some cases, that increase in rendering speed can be much larger. The RT Cores accelerate ray casting – the process of intersecting a light ray with all of the objects in a scene – so scenes with complex geometry but relatively simple shading, such as large massing models, benefit greatly, as do those with heavy instancing.

"OptiX is really good at heavily instanced scenes, such as those with a lot of grass," says Miller. "In those cases, we see a really big speed bump, because we are getting the benefits of the RT Cores and then OptiX as well."

RTX acceleration also lies at the heart of Project Lavina, Chaos Group's upcoming product for manipulating V-Ray scenes in real time. The technology, currently in closed beta, enables users to explore scenes in a completely ray traced environment, iterating on designs and choosing camera positions.

"The challenge with traditional renderers is that they were designed to run on processors in a certain way. That can change dramatically when suddenly rays are free," says Miller. "Project Lavina was designed with the assumption that the user would have an RTX card, so we're gleaning as much performance from the RT Cores as is possible. With V-Ray GPU, we have some scenes that can go 100% faster [on RTX cards]. With Lavina, it's more like 400-500% faster."

An image from Chaos Group's 2019 technology preview of Project Lavina. Scene courtesy of Evan Butler.


Proxies in use with V-Ray GPU on RTX

MORE COMPLEX, MORE BEAUTIFUL REAL-TIME RENDERS

One firm currently trialing Project Lavina is KPF (Kohn Pedersen Fox Associates). The international architecture practice works on projects ranging from small commercial units to the largest urban developments, including airports and supertall structures: the firm has worked on five of the world's ten tallest towers.

In his tests, Paul Renner, KPF's Visualization Manager and BIM Practice Specialist, trialed Project Lavina on assets of a similar complexity, including a 31.5-million-polygon massing model of New York City.

"I was blown away," he comments. "[The scene] runs smoothly in Lavina on a NVIDIA Quadro RTX 6000. Being able to add in HDRIs and give it different environment moods [in real time] was really magical."

Architectural practice KPF uses a range of RTX-accelerated tools, including V-Ray GPU and Enscape, in its work.
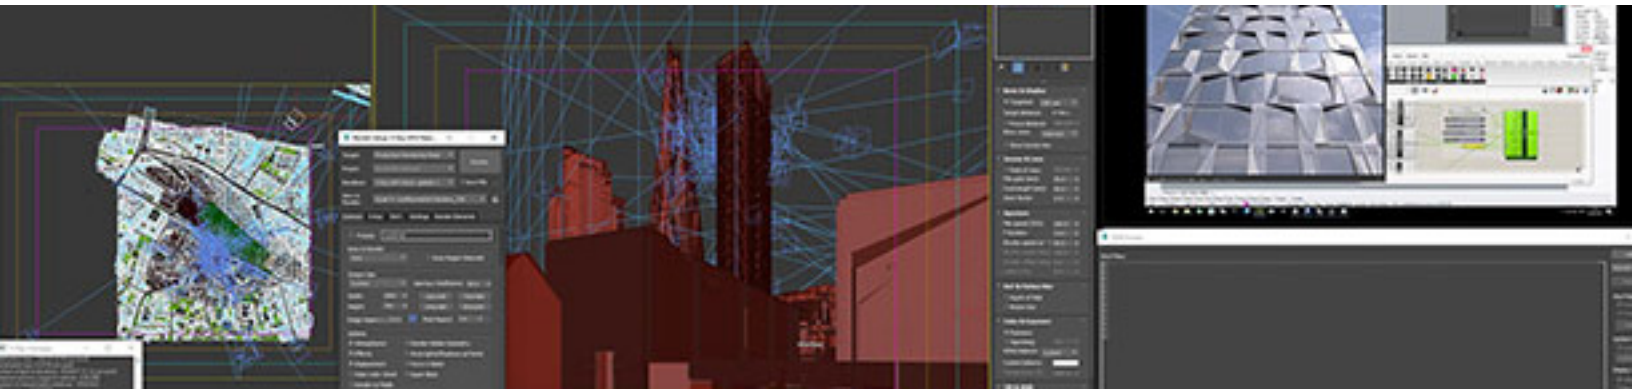
Paul Renner's 31.5-million-polygon massing model of New York City, rendered from within Project Lavina.
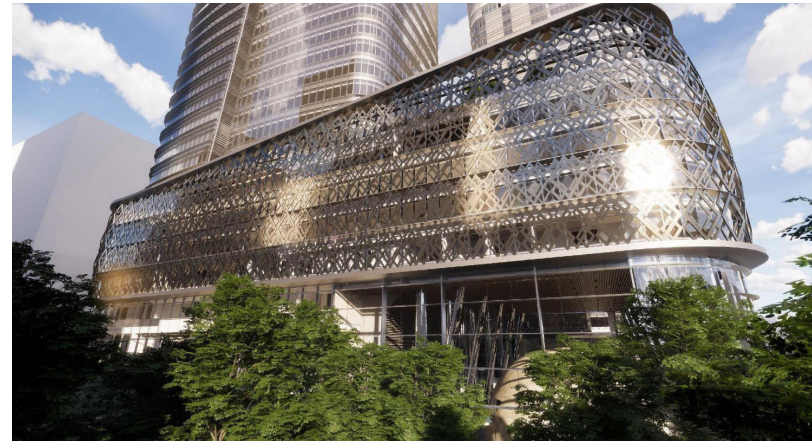
An image generated in Enscape by KPF. RTX acceleration improves the visual quality of glass in such renders.

KPF is currently installing NVIDIA Quadro RTX 6000 nodes in its in-house render farm, as well as rolling out NVIDIA Quadro RTX 4000 cards in the workstations of its 700-strong design team, where they help to accelerate RTX-capable tools like V-Ray and Enscape, the real-time rendering and virtual reality plugin for Revit, SketchUp, Rhino and ArchiCAD.

"We encourage the team to learn [visualization] skills, because they can do a better job of designing that way," says Renner.

"With V-Ray GPU and Enscape, we can do as much in-house as possible: not just massing or basic-shaded models. Now a designer can produce really beautiful images faster than ever."

Thanks to the NVIDIA Quadro RTX cards, the visual quality of the images that KPF can generate in Enscape has increased, particularly when it comes to processor-intensive reflections and refractions in glass curtain walls.

RTX acceleration also helps KPF to produce more ambitious animations of its projects. "On RTX, it takes Enscape 15 minutes to output a 4K animation a minute or two long, whereas before, it might have taken three hours," says Renner. "It's going to change how we present our work. We don't have to budget time to do a huge animation: we can just throw one in every time, and that helps to explain the design principles of the project in a better way."

NEW INTERACTIVE AND AI-DRIVEN VIRTUAL REALITY WORKFLOWS

NVIDIA Quadro RTX cards don't just help studios to generate pre-rendered content. The hardware also powers the new DXR ray tracing pipeline in Unreal Engine, making it possible to create immersive virtual reality projects.

One company taking advantage of the transformative power of game engines in architectural work is Zoan. The Finland-based firm began developing its first fully VR project – an interactive representation of the city of Helsinki – in 2016, and is currently transitioning into a dedicated virtual reality studio.

In 2019, Zoan won the new Unreal Studio DXR Ray Tracing category at the CGarchitect Architectural 3D Awards, turning an existing client model of a series of office spaces into a fully interactive environment. Although the firm has yet to use DXR ray tracing on a commercial VR project, it is already taking advantage of the power of RTX hardware in its work, using a mixture of NVIDIA GPUs in production, including the NVIDIA Quadro RTX 6000.

Zoan's winning entry in the Unreal Studio DXR Ray Tracing category at the CGarchitect Architectural 3D Awards.

Zoan's first major VR project was Virtual Helsinki, an interactive recreation of Finland's capital city.

RTX technology accelerates Zoan's entire pipeline, beginning with making initial design decisions interactively in VR with a client. "RTX gives you a tool for testing," says Alfonso Escalona Penott, the studio's Lead Technical Artist for Unreal Engine. "The client can see for themselves if they are committed to an idea, or if they want changes."

The NVIDIA Quadro RTX cards also increase the fidelity of the imagery that Zoan generates. "We're trying to deliver our final renders with real-time reflections using RTX acceleration to produce very large, high-quality images, and for that, we

need very powerful Quadro cards," says Penott. "[With RTX] previsualisation is also faster and more precise. Having a faster preview really helps workflow. We can create bigger projects in the time [available]."

Zoan is now experimenting with new ways to make BIM data interactive, enabling changes to a BIM file to generate updates to a virtual reality project in real time. "One of our goals is to use AI to generate the decorations for a room," says Penott. "If the architect modifies its layout, [the AI] will repopulate the scene automatically."

> ## "We're trying to deliver our final renders with real-time reflections using RTX acceleration to produce very large, high-quality images, and for that, we need very powerful Quadro cards,"
>
> *Alfonso Escalona Penott*
> *Zoan, Lead Technical Artist for Unreal Engine*

NVIDIA QUADRO RTX GRAPHICS CARDS: TRANSFORMING THE WAY THAT ARCHITECTS WORK

Through their Turing architecture, NVIDIA Quadro RTX GPUs are changing the way that architectural professionals work. Their industry-leading compute performance and memory capacity makes them ideal for demanding design and visualization projects, while their dedicated Tensor and RT Cores open up new AI-driven and real-time workflows.

For studios like KPF, these new workflows make day-to-day work not merely faster, but more creative. "Right now, if we have a week between client reviews, we have to figure out how much time we have for rendering and animation as well as the modeling and design," says Paul Renner. "With [RTX acceleration], we can produce final imagery as we go. It changes our work from a two-step process of design and wait, to one of just design."

For additional information on how NVIDIA Quadro RTX technology is being used to transform architectural visualization, visit **WWW.PNY.COM/AEC.**

PNY.