
A E S T R U C T

Neural Rendering: The Next Frontier

What It Is, Why It Matters & What It Means
for VFX & Animation Studios

February 2026

C O N F I D E N T I A L

Executive Summary

In January 2026, NVIDIA CEO Jensen Huang stood on stage at CES and said five words that should matter to every VFX and animation studio: “The future is neural rendering.”

He wasn’t being vague. He was describing a specific, measurable shift in how images will be created inside computers. Traditional rendering — the process that turns 3D scenes into the images you see on screen — works by simulating the physics of light. It’s powerful, it’s proven, and it’s expensive. Neural rendering replaces parts of that physics simulation with trained AI models that produce equivalent (or better) results in a fraction of the time.

This isn’t science fiction. It’s shipping in production tools today. NVIDIA’s OptiX 9.0, Unreal Engine 5, V-Ray 7, and Nuke 17 all include neural rendering features. The 2025 production of Superman used 3D Gaussian splatting — a core neural rendering technique — making it the first major motion picture to deploy the technology. Microsoft’s RenderFormer research demonstrated a complete neural rendering pipeline that’s 40–120x faster than traditional path tracing for certain scenes.

This report explains what neural rendering actually is (without requiring a computer science degree), maps what’s available today versus what’s coming in the next 12–24 months, and lays out the practical opportunities, cost implications, and risks for studios evaluating this technology.

The Short Version

Neural rendering uses AI to do work that used to require brute-force computation. It’s already cutting render times by 60–80% through denoising and upscaling. The next wave — neural textures, neural lighting, and eventually full neural scene rendering — will reshape how studios budget time, hardware, and talent. Studios that understand this shift now will have a significant head start.

1. What Is Neural Rendering?

1.1 The Problem It Solves

To understand neural rendering, you need to understand what traditional rendering does and why it's so resource-intensive.

When a VFX or animation studio creates a 3D scene — characters, environments, lighting — the computer needs to turn that 3D data into a 2D image for the screen. This process is called rendering. Traditional rendering does this by simulating light physics: tracing millions of light rays as they bounce between surfaces, calculating how each ray interacts with different materials (glass, skin, metal, fabric), and accumulating all that information into a final image.

More rays equals cleaner images. But more rays also equals more time and more hardware. A single frame of a Pixar film can take 24 hours to render on powerful hardware. An episode of a VFX-heavy streaming show might require thousands of compute-hours across a render farm. This is the fundamental bottleneck: traditional rendering trades time and money for image quality.

1.2 The Core Idea

Neural rendering asks a different question: instead of simulating every light ray from scratch, what if we trained an AI model to predict what the final image should look like?

The AI learns by studying millions of examples — pairs of “rough” and “finished” images — until it can reliably predict the clean result from incomplete information. Think of it like an expert painter who's seen so many landscapes that they can sketch a photorealistic scene from a rough thumbnail. The physics knowledge is baked into the model's training, not recalculated from scratch every time.

This is fundamentally different from traditional rendering. Instead of more computation (more rays, more time, more hardware), neural rendering uses pattern recognition to achieve equivalent results with dramatically less work.

1.3 A Simple Analogy

Traditional vs. Neural Rendering

Traditional rendering: You want to know the temperature in your city. You build a complete atmospheric simulation modeling every air molecule and weather pattern from first principles. It takes hours and massive computational resources, but it's physically accurate. Neural rendering: You train an AI on decades of weather data for your city. It learns the patterns and predicts the temperature in seconds with 97% accuracy. It didn't simulate the physics — it learned the answer. Both approaches give you the temperature. One costs dramatically less to run.

1.4 The Spectrum of Neural Rendering

Neural rendering isn't one technology. It's a family of techniques that range from "modest AI assist" to "complete replacement of the rendering pipeline." Here's the landscape, from least to most ambitious:

Approach	What It Does	Where It Is Today
AI Denoising	Cleans up noisy, partially-rendered images. Renders faster because fewer light rays are needed.	Standard. Shipping in all major render engines. 92% of top studios use it.
AI Upscaling (DLSS)	Renders at lower resolution, AI reconstructs full resolution. DLSS 4.5 generates up to 15 of every 16 pixels with AI.	Mature. DLSS 4.5 shipping January 2026. Used in real-time and offline workflows.
Neural Textures	AI-compressed textures that adapt to viewing angle and lighting. 8x memory savings vs. traditional textures.	Early production. NVIDIA RTX Kit, shipping with Blackwell GPUs (2025–2026).
Neural Radiance Cache	AI learns multi-bounce indirect lighting during the render, inferring infinite bounces from 1–2 traced bounces.	Production-ready. Shipping in RTX Remix, RTXGI SDK. Used in Half-Life 2 RTX.
Neural Materials	AI compresses complex multi-layer shaders (porcelain, silk) for real-time use. 5x faster material processing.	Early production. Announced in NVIDIA RTX Kit (2025).
3D Gaussian Splatting	Represents 3D scenes as clouds of colored points instead of geometry. Real-time, photorealistic scene capture.	Production use. Used in Superman (2025). Nuke 17, OctaneRender 2026, Unreal Engine support.
Neural Radiance Fields (NeRF)	Trains a neural network to represent a complete 3D scene from photographs. Novel viewpoints generated on demand.	Research/early production. Integrated into AI video generators and virtual production tools.
Full Neural Pipeline (RenderFormer)	Replaces the entire rendering pipeline with a transformer model. No traditional graphics computation.	Research only. Microsoft SIGGRAPH 2025 paper. 40–120x faster but limited to simple scenes.

The key insight: we're not waiting for some future breakthrough. Studios are already using the first four or five items on this list. The question is how fast the bottom of the list moves into production.

2. What's Shipping Today

Let's separate what's real and production-proven from what's promising but unproven. This section covers only technologies that are in commercial products and being used on actual productions right now.

2.1 NVIDIA's Neural Rendering Stack

NVIDIA is the dominant force here. Their hardware (RTX GPUs with Tensor Cores) and software (OptiX, DLSS, RTX Kit) form the foundation that most neural rendering runs on.

OptiX 9.0 and Cooperative Vectors

OptiX is NVIDIA's ray tracing engine — the software that most major render engines use under the hood for GPU rendering. Version 9.0 introduced “cooperative vectors,” which is a technical way of saying: small AI neural networks can now run directly inside the ray tracing process itself, accelerated by dedicated AI hardware (Tensor Cores) on RTX GPUs.

Why this matters for studios: it means AI inference and traditional rendering can happen simultaneously on the same GPU, in the same shader program. You don't need a separate AI step — the neural network is woven into the rendering pipeline. This is the infrastructure that enables neural textures, neural materials, and neural radiance caching.

Critically, cooperative vectors are also being added to Microsoft's DirectX (Shader Model 6.9) and the Vulkan API. This is becoming an industry standard, not just an NVIDIA feature.

DLSS 4.5

DLSS (Deep Learning Super Sampling) is NVIDIA's AI upscaling technology, now in its fourth major generation. DLSS 4.5, announced at CES 2026, uses a second-generation transformer model for frame generation. NVIDIA's own claim: up to 15 of every 16 pixels on screen can be AI-generated, with the GPU only rendering 1 in 16 at native resolution.

For studios, the production implication is straightforward: render at a fraction of the resolution, let AI reconstruct the rest. This directly translates to faster interactive previews, faster final renders, and lower hardware requirements for real-time work (virtual production, previz, interactive reviews).

RTX Kit: The Full Neural Toolkit

NVIDIA's RTX Kit bundles their neural rendering technologies into a developer-facing suite. The pieces that matter most for VFX and animation:

- **Neural Texture Compression (NTC):** Compresses textures using AI, reducing VRAM usage by up to 8x. In NVIDIA's demo, textures that consumed 520MB with traditional compression occupied just 80MB with NTC. For studios working with massive texture sets (film environments, open-world games), this is a direct capacity multiplier.

- **Neural Materials:** Compresses complex multi-layer shaders — the kind used for materials like porcelain, silk, or wet skin — so they can run at real-time frame rates. Material processing is up to 5x faster. This bridges the gap between “offline” film-quality shaders and “real-time” game shaders.
- **Neural Radiance Cache (NRC):** AI learns how indirect light bounces through a scene, then predicts the result for additional bounces without actually tracing those rays. The network trains during the render session itself, adapting to the specific scene. This ships in RTX Remix and the RTXGI SDK, and was used in Half-Life 2 RTX.
- **RTX Mega Geometry:** Handles scenes with extremely dense geometry — billions of triangles — in real time. Debuting in Alan Wake 2. For VFX, this means environments that would previously choke real-time engines can now be interactive.

2.2 3D Gaussian Splatting in Production

3D Gaussian splatting (3DGS) deserves special attention because it’s already in production at the highest level.

Instead of representing a 3D scene as a mesh of triangles (the standard approach since the 1970s), Gaussian splatting represents scenes as clouds of millions of tiny colored, semi-transparent blobs. Each blob has a position, color, opacity, and shape. Rendered together, they produce photorealistic images of real-world scenes captured from photographs or video.

Why Studios Care

- **Speed:** Capture a real location with a camera, process the footage, and you have a photorealistic 3D environment in hours instead of weeks of manual modeling.
- **Quality:** Gaussian splats capture real-world lighting, materials, and geometry with a fidelity that’s extremely difficult to match by hand.
- **Real-time:** Unlike NeRFs (which are slow to render), Gaussian splatting runs in real time, making it suitable for virtual production on LED volumes.

Production Milestones

- **Superman (2025):** First major motion picture to use dynamic Gaussian splatting in its VFX pipeline. Houdini-based 3DGS tools were integrated into professional workflows.
- **Nuke 17 (2025):** The Foundry announced native 3DGS support in the open beta of their industry-standard compositing software.
- **OctaneRender 2026 (OTOY):** Full path-traced Gaussian splats, with 2027 adding splats as a native export from synthetic scenes.
- **Volinga Plugin Pro:** Integrates 3DGS into Unreal Engine with relighting, HDRI, and ACES color management — the features needed for professional production.
- **V-Ray 7 (Chaos):** Added 3D Gaussian splatting on GPU as part of the V-Ray 7 update.

Industry Signal

Early adopters report up to 60% reduction in VFX iteration time when using Gaussian splatting for environment capture and virtual production. The technology's arrival in Nuke, V-Ray, OctaneRender, and Unreal Engine means it's no longer a standalone research tool — it integrates into existing pipelines.

2.3 Render Engine AI Features: Current State

Engine	Neural Rendering Features (2025–2026)
V-Ray 7 (Chaos)	NVIDIA + Intel AI denoising, AI Enhancer (diffusion-based), 3D Gaussian splatting on GPU, per-device GPU control for simultaneous render + denoise
Arnold 7.4.x (Autodesk)	NVIDIA OptiX denoiser integration, continued use as primary renderer for film VFX at major studios
Redshift (Maxon)	AI denoiser, fast GPU rendering for motion graphics and animation under tight deadlines
Karma (SideFX / Houdini)	AI denoising, deep integration with Houdini simulation tools, Gaussian splatting workflows via SOP context
Unreal Engine 5 (Epic)	DLSS 4.5, Lumen (dynamic GI), Nanite (virtualized geometry), RTX Mega Geometry, Neural Radiance Cache, neural shader support
OctaneRender 2026 (OTOY)	Path-traced Gaussian splats, AI-accelerated rendering, ORBX scene format for neural assets
Nuke 17 (Foundry)	Native 3D Gaussian splat support, CopyCat ML node, AI-assisted compositing and rotoscoping

3. What's Coming Next

The line between “shipping today” and “coming soon” is thin and moving fast. Here's what the next 12–24 months look like based on announced roadmaps and credible research.

3.1 Microsoft RenderFormer: The Full Replacement Vision

In mid-2025, Microsoft Research published a paper (presented at SIGGRAPH 2025) that demonstrated something no one had done before: a neural network that can perform complete 3D rendering — including global illumination — with zero traditional graphics computation.

RenderFormer uses a transformer architecture (the same type of AI model that powers ChatGPT and other large language models) to convert a 3D scene description directly into a rendered image. No ray tracing. No rasterization. Just a neural network that's learned what light does.

The Numbers

- **Speed:** 40–120x faster than traditional path tracing (Blender Cycles) for comparable scenes.
- **Architecture:** 205 million parameters. Two-stage transformer: one for view-independent effects (shadows, diffuse light), one for view-dependent effects (reflections, specular highlights).
- **Training:** Trained on Objaverse (800,000+ 3D objects), no per-scene fine-tuning required.
- **Current limitation:** Works on scenes up to ~7,000 triangles at 512×512 resolution. Production VFX scenes contain millions to billions of triangles. This is a research proof of concept, not a production tool — yet.

Why This Matters Even Though It's Early

RenderFormer is significant because it proves the concept works end-to-end. A neural network can learn the entire rendering pipeline — geometry processing, light transport, material interaction, global illumination — without any hand-coded graphics algorithms. The 40–120x speed advantage suggests that when this approach scales to production complexity, the economics of rendering will change fundamentally. Microsoft has open-sourced the code on GitHub, which means the research community is already building on it.

3.2 NVIDIA's Stated Direction

Jensen Huang's CES 2026 comments were unusually direct. He described neural rendering as “basically DLSS” applied more broadly — the idea that AI generates most of the pixels, and traditional rendering provides the minimum required “ground truth” for the AI to work from.

He also suggested the RTX 5090 may represent “the pinnacle of traditional raster” — meaning NVIDIA doesn’t plan to keep scaling traditional GPU shader power indefinitely. Instead, future performance gains will come from better AI inference, not bigger shader arrays. For the first time in five years, NVIDIA did not announce a new GPU at CES.

The practical implication: studios that invest in RTX-class hardware are buying into a platform that will become more capable through software and AI model updates, not just hardware refreshes. The Tensor Cores that accelerate AI inference today will be the primary performance engine of tomorrow.

3.3 4D Gaussian Splatting

Static Gaussian splatting captures a frozen moment. 4D Gaussian splatting adds time — capturing dynamic scenes with motion, performance, and changing lighting. Companies like Gracia (which received \$1.7 million in funding for Hollywood-scale 4DGS tools) are solving the hardest problems: handling fast complex motion, avoiding temporal artifacts, and supporting slow-motion playback.

For VFX and animation studios, 4DGS means capturing real performances in full volumetric detail — an actor’s performance, a practical explosion, a choreographed stunt — and bringing that captured reality into the digital pipeline as a manipulable 3D asset, not just a 2D plate.

3.4 Neural Rendering in AI Video Generation

The AI video generation tools (Runway, Sora, Kling) are starting to integrate neural radiance fields and Gaussian splatting internally. This means AI-generated video is gaining spatial consistency — understanding 3D geometry rather than just generating plausible-looking frames. For VFX pipelines, this opens the door to using generative AI for background plates, environment extensions, and previz that maintains correct 3D perspective and lighting.

4. The Business Case

Technology analysis without cost implications is just trivia. Here's what neural rendering means for studio economics.

4.1 Direct Cost Savings

Neural Rendering Feature	Where Savings Come From	Estimated Impact
AI Denoising	Fewer ray samples needed per frame. Same quality at 1/5 to 1/3 of the render time. Directly reduces cloud compute costs or frees on-prem nodes.	60–80% render time reduction
DLSS / AI Upscaling	Render at lower resolution (1/4 to 1/16 of target), AI reconstructs full res. Massive compute savings for real-time and interactive workflows.	Up to 4x performance gain
Neural Texture Compression	8x VRAM reduction means larger scenes fit on fewer (or cheaper) GPUs. Reduces memory-related render failures.	Up to 8x VRAM savings
Neural Materials	Film-quality materials run at game frame rates. Eliminates the need for separate "real-time proxy" and "offline hero" material versions.	5x material processing speedup
Gaussian Splatting (Environment Capture)	Days of manual environment modeling replaced with camera capture + processing. Photorealistic real-world environments in hours.	60–90% environment creation time reduction
Neural Radiance Cache	AI handles multi-bounce lighting. Fewer rays traced for equivalent lighting quality. Significant speedup for interior and indirect-lit scenes.	2–4x lighting computation reduction

4.2 Hardware Investment Implications

The shift to neural rendering has a direct impact on hardware purchasing decisions:

- **Tensor Cores are no longer optional.** Every neural rendering feature requires NVIDIA RTX GPUs with Tensor Cores. Studios running older pre-RTX hardware cannot access any of these capabilities. The ROI on RTX upgrades is now clear and immediate.
- **VRAM matters more than raw compute.** Neural textures and Gaussian splatting are memory-intensive. Studios should prioritize GPUs with more VRAM (RTX 4090 at 24GB, or upcoming RTX 5090) over raw shader performance. This is a shift from traditional rendering where compute speed was the primary bottleneck.

- **The render farm gets smaller (or cloudier).** If AI denoising cuts render times by 60–80%, a studio needs fewer on-prem render nodes for the same throughput. The savings can fund cloud burst capacity for peak loads or be reinvested in other production areas.

4.3 Revenue and Competitive Implications

The studios that will benefit most from neural rendering aren't the ones that use it to cut costs. They're the ones that use it to take on work they couldn't handle before:

- **Bid more competitively:** Lower per-frame render costs mean tighter bids without sacrificing margins.
- **Deliver more ambitious work:** The same budget buys more VFX shots, higher quality, or both.
- **Faster turnaround:** Interactive neural rendering (previews, client reviews, direction changes) compresses production timelines.
- **New service offerings:** Gaussian splatting-based environment scanning, AI-accelerated previz, neural lighting for virtual production — these are new revenue streams for studios that build the capability.

The Reinvestment Pattern

Total VFX budgets don't always decrease when studios adopt neural rendering. What typically happens is producers reinvest the efficiency savings into more ambitious visual work. A show that could afford 200 VFX shots with a traditional pipeline can now deliver 300 shots at the same budget. This is good news: margins are protected while quality rises.

5. Risks & Limitations

Neural rendering is not a magic wand. Here's an honest assessment of the risks, organized by how likely they are to affect a studio in the near term.

5.1 Technical Risks

Quality Control Challenges

AI denoising can smooth over fine geometric detail or introduce subtle artifacts — particularly in areas with complex specular highlights, hair, fur, or translucent materials. These artifacts are often invisible at interactive resolution but visible on a cinema screen or in a client's 4K QC pass. Studios need to build AI artifact checks into their QC pipeline, not assume the AI output is always correct.

Gaussian Splatting Limitations

Gaussian splats are extraordinary for static and slowly-changing environments. But they struggle with fast complex motion, thin structures (fences, wires, hair), and transparent or highly reflective surfaces. 4D Gaussian splatting addresses some of these issues but is still maturing. For hero character work with close-up detail, traditional geometry-based approaches remain necessary.

The “Black Box” Problem

When a traditional render looks wrong, an artist can trace the problem to a specific light, material, or geometry setting and fix it. When a neural render looks wrong, the cause is inside a trained model with millions of parameters. Debugging is harder, and fixes may require retraining or switching models rather than adjusting a slider. This is a real workflow concern for studios that need tight control over final pixel output.

5.2 Business & Operational Risks

Vendor Concentration

NVIDIA dominates the neural rendering hardware and software ecosystem. Cooperative vectors, OptiX, DLSS, and RTX Kit are all NVIDIA technologies. While Microsoft is adding cooperative vector support to DirectX, and AMD's FSR Redstone is an alternative AI upscaler, the practical reality is that most neural rendering capabilities require NVIDIA hardware. Studios should be aware of this dependency.

Training Data and IP Concerns

Some AI models used in neural rendering were trained on datasets scraped from the internet (LAION-5B, for example). This creates a legal gray area around whether generated outputs are derivative works. For studios working on IP-sensitive projects — particularly with major rights

holders — understanding what's in the training data of their AI tools is a due diligence requirement, not a nice-to-have.

Workforce Transition

An estimated 21% of film, television, and animation jobs face AI-related consolidation by 2026. Neural rendering specifically affects lighting TDs, texture artists, and environment artists — roles where AI can automate significant portions of the work. The studios handling this well are treating AI as a capability amplifier for existing talent, not a replacement strategy. The ones handling it poorly are accumulating labor risk.

Aesthetic Convergence

Over-reliance on AI generation risks a convergence in visual style. When many studios use the same AI models trained on the same data, outputs start to share a recognizable “AI look.” Audiences are increasingly able to spot AI-generated imagery. Studios that rely on AI for speed but maintain strong art direction and human creative oversight will differentiate; studios that use AI as a style crutch will produce interchangeable work.

5.3 Standards and Interoperability

The neural rendering ecosystem is fragmented. NVIDIA's formats, OpenUSD, SMPTE metadata schemas, and vendor-specific implementations don't always play well together. Gaussian splatting alone has multiple competing file formats. Until standards mature — and the OpenUSD effort is making progress — studios face engineering overhead getting neural rendering assets to move cleanly between tools and departments.

6. Maturity Assessment

Not all neural rendering technologies are at the same stage. This matrix maps readiness levels against production impact to help prioritize where to invest attention.

Technology	Maturity	Production Impact	Recommendation
AI Denoising	Mature — Standard	High: 60–80% render time savings	Adopt now. If you're not using this, you're overspending on rendering.
AI Upscaling (DLSS 4.5)	Mature — Production	High: Up to 4x performance for real-time workflows	Adopt now for previz, interactive review, and virtual production.
3D Gaussian Splatting	Early Production	High: 60–90% faster environment creation	Evaluate now. Build test workflows. Production use for non-hero environments.
Neural Texture Compression	Early Production	Medium: 8x VRAM savings	Adopt when available in your pipeline tools. Low risk, high upside.
Neural Radiance Cache	Early Production	Medium: 2–4x lighting speedup	Evaluate for real-time and virtual production workflows.
Neural Materials	Pre-Production	Medium: 5x material speedup	Watch. Will matter when film-quality materials need to run in real time.
4D Gaussian Splatting	Emerging	Potentially High: Volumetric performance capture	Monitor. Maturing fast. Evaluate for performance-capture-heavy projects.
Full Neural Pipeline (RenderFormer)	Research	Transformative (long-term)	Monitor. Not ready for production, but directionally important.

7. Practical Recommendations

Here's how we'd suggest a VFX or animation studio approach neural rendering, ordered by urgency and impact:

Immediate Actions (This Quarter)

1. **Audit your AI denoising usage.** If you're not using AI denoising in every render engine in your pipeline, turn it on and benchmark the difference. This is the lowest-risk, highest-return action available. V-Ray 7, Arnold, Redshift, and Karma all support it.
2. **Evaluate your GPU fleet.** Make an inventory. Any GPU without Tensor Cores is locked out of neural rendering. Plan your upgrade path. Prioritize VRAM (24GB+) over raw shader count.
3. **Test Gaussian splatting on a real project.** Pick a non-critical environment or establishing shot. Capture reference photography, process it through a Gaussian splatting pipeline, and evaluate the result against your quality standards. The goal is learning, not shipping.

Near-Term Investments (Next 6–12 Months)

4. **Build hybrid cloud burst capability.** If you haven't already, set up cloud rendering (AWS Deadline Cloud or equivalent) so you can scale beyond your on-prem farm. AI denoising + cloud burst is a powerful combination that fundamentally changes your cost structure.
5. **Integrate AI compositing tools.** Mocha Pro 2026, Nuke CopyCat, and After Effects' AI roto tools save enormous artist-hours. These aren't neural rendering per se, but they're part of the same AI-accelerated pipeline.
6. **Train your team.** The biggest barrier isn't technology — it's mindset. Artists and supervisors need hands-on experience with AI tools to understand what they're good at and where human judgment is still essential. Budget for learning time.

Strategic Positioning (12–24 Months)

7. **Plan for neural rendering as default infrastructure.** Within two years, AI denoising, neural textures, and neural lighting will be standard features in all major render engines. Build your pipeline and hiring plans around this assumption.
8. **Watch the standards landscape.** OpenUSD, SMPTE metadata standards, and cooperative vector APIs are converging. Don't commit to proprietary formats that may not survive standardization.
9. **Consider new service offerings.** Gaussian splatting-based environment scanning, AI-accelerated virtual production, neural lighting for LED volumes — these are emerging revenue opportunities for studios that build the capability early.

8. Where This Goes

Based on what's announced, what's shipping, and where the research is pointing:

2026: Neural Rendering Becomes Table Stakes

- AI denoising is fully standard across all render engines. Studios not using it are operating at a measurable cost disadvantage.
- DLSS 4.5 and neural upscaling are the default for interactive and real-time workflows. Rendering at native resolution becomes optional for everything except final hero frames.
- 3D Gaussian splatting moves from early adoption to mainstream, with native support in Nuke, V-Ray, OctaneRender, and Unreal Engine.
- Virtual production LED volumes adopt neural radiance caching for real-time global illumination, replacing baked lighting approaches.

2027–2028: The Pipeline Shifts

- Neural textures and neural materials become standard, collapsing the gap between offline and real-time rendering quality. “Proxy” and “hero” versions of assets converge into a single asset that works everywhere.
- 4D Gaussian splatting matures for performance capture, offering a viable alternative to traditional motion capture for many use cases.
- AI-driven render farm scheduling finally ships as a standard feature, using historical data to predict render times and optimize resource allocation automatically.
- Full neural rendering (RenderFormer-style) handles non-hero elements: distant backgrounds, crowd shots, environment extensions. Traditional ray tracing remains for hero close-ups.

The Bigger Picture

Jensen Huang's statement that the RTX 5090 may be “the pinnacle of traditional raster” isn't a throwaway line. It signals that the entire GPU industry — NVIDIA, AMD, Microsoft — is pivoting toward neural computation as the primary path for visual quality improvement. The studios that understand this shift and build for it will have fundamentally different cost structures and capabilities than those that don't.

Neural rendering doesn't make artists obsolete. It changes what artists spend their time on. Instead of waiting for renders and building assets from scratch, they're directing AI tools, refining AI output, and making creative decisions faster. The studios that thrive will be the ones that use neural rendering to amplify their artists' judgment, not replace it.

This analysis was prepared by AEstruct, Inc. for prospective client review. All data points are sourced from publicly available vendor documentation, trade publications, academic papers, and industry reports as of February 2026.

Questions? Reach out to discuss how these findings apply to your specific production environment and technology roadmap.

A E S T R U C T
Create More. Manage Less.

© 2026 AEstruct, Inc. All Rights Reserved.