Live Webinar:

NVIDIA Studio Tools Roundup

GPU ACCELERATION AND AI
ENHANCED VISUAL APPLICATIONS

PNY NVIDIA

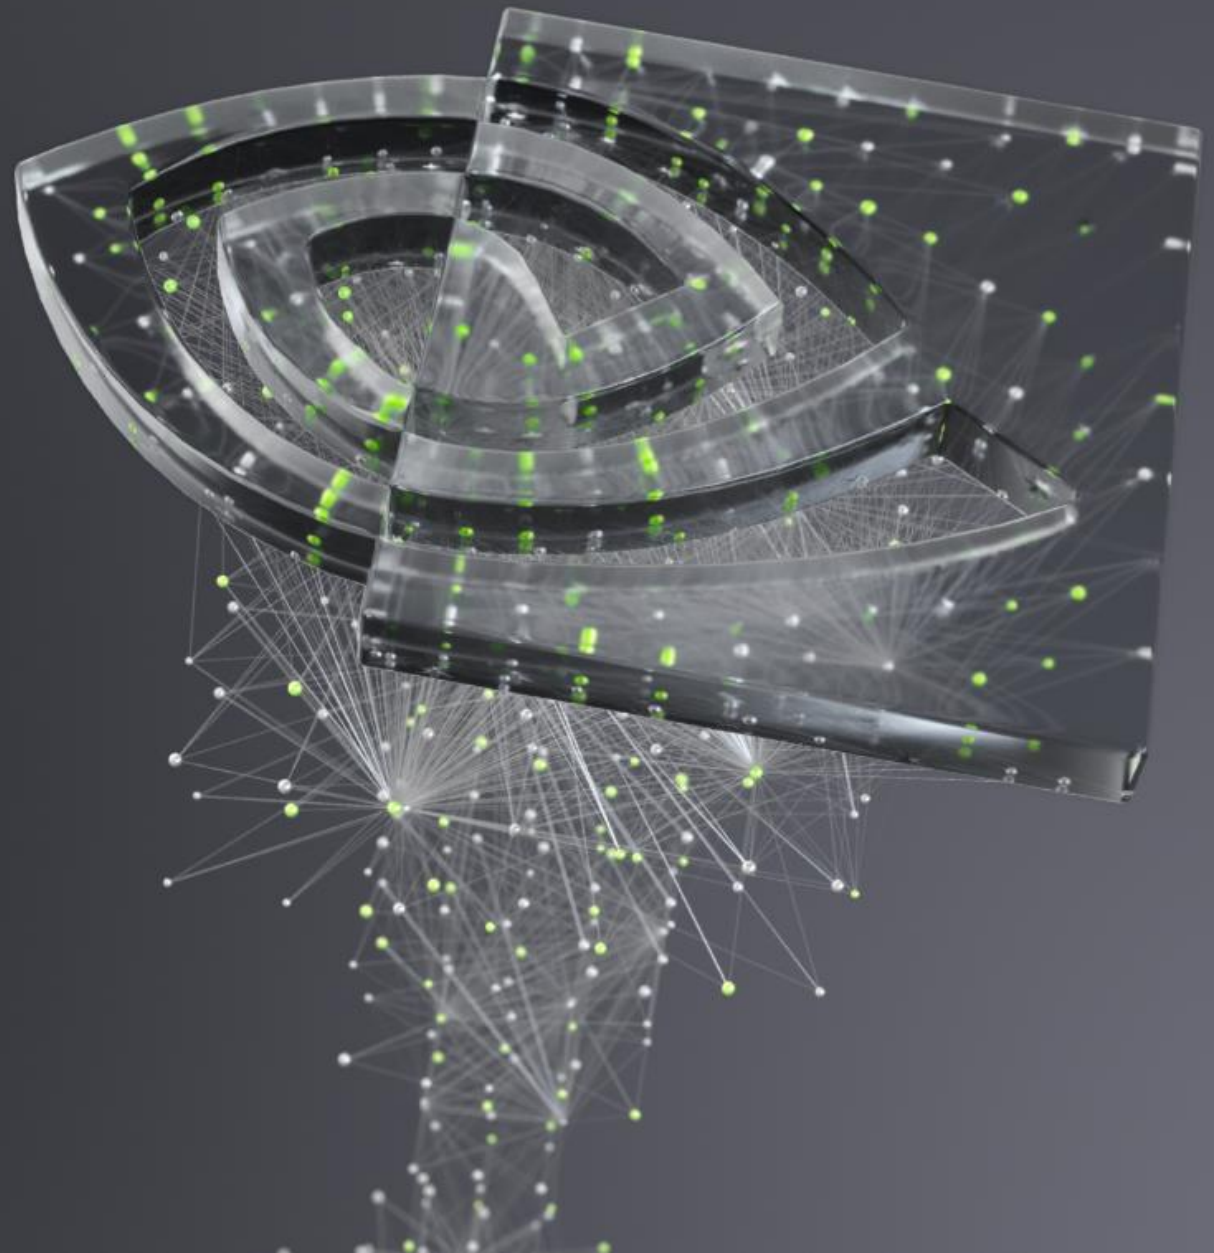

ENTERPRISE MEDIA SOLUTIONS

Rick Grandy – Principal Solutions Architect - NVIDIA

Expanding the Universe of Visual Computing

- Workloads growing in size and complexity
- Advanced technologies such as AI, data analytics, visualization becoming more common
- Need to maximize perf / watt / foot^2

RTX ON

Performance and Realism in the Creative Process Unlike Ever Before

- Lighting-fast productivity

Design and create in real-time

- Large-scale models and scenes

96 GB on NVIDIA Quadro RTX

- Real-time photoreal ray tracing

New levels of interactive realism

- AI-augmented graphics

More efficient workflows, enhanced image quality

NVIDIA RTX ACCELERATED APPS

50+ Apps Available Now

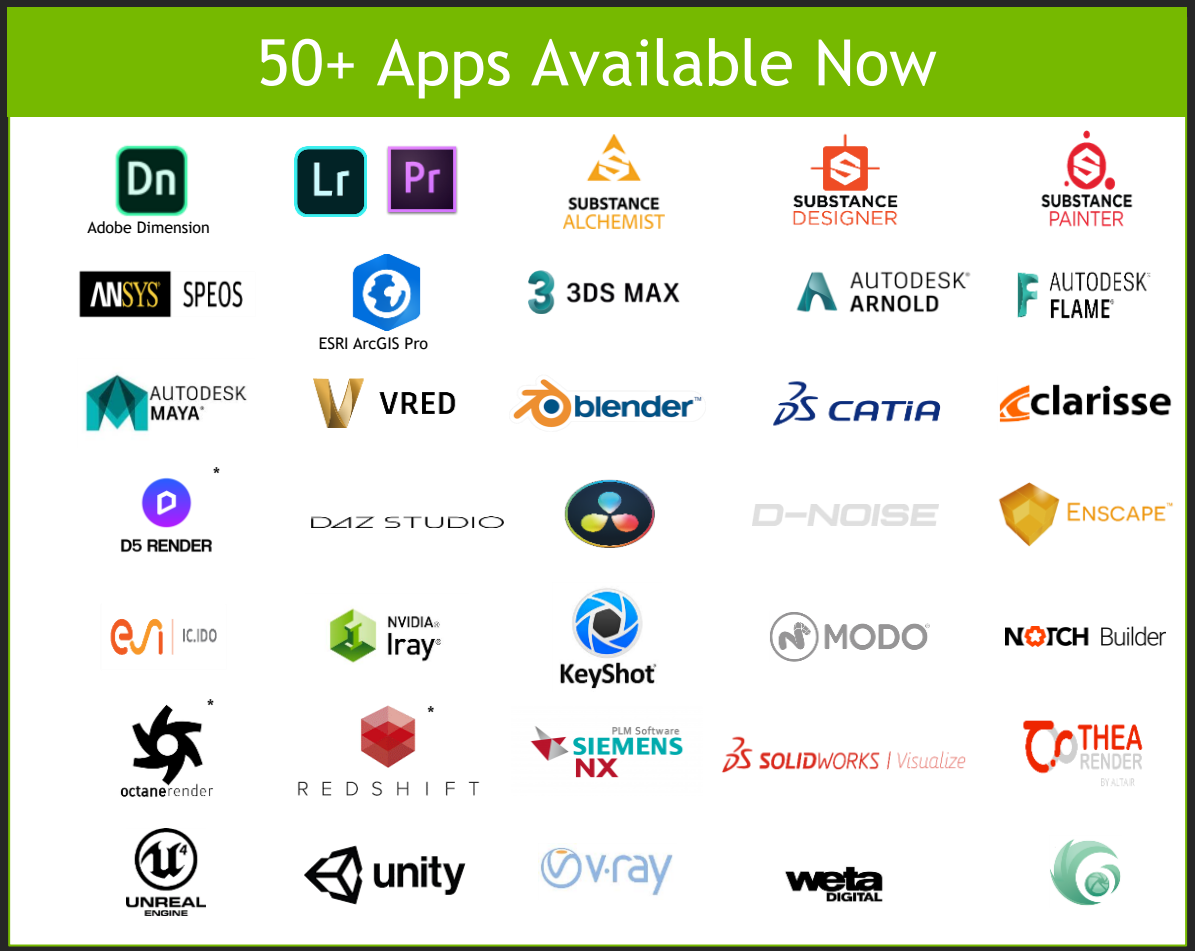

*Marked Applications in Beta
*Applications support some or all RTX features

Adobe

RTX-accelerated features for millions of 2D/3D artists and graphic designers

Dimension

Create environments in photorealistic 3D, featuring RTX ray tracing

Substance Alchemist

Create and blend materials with ease, featuring RTX-accelerated AI

Substance Designer

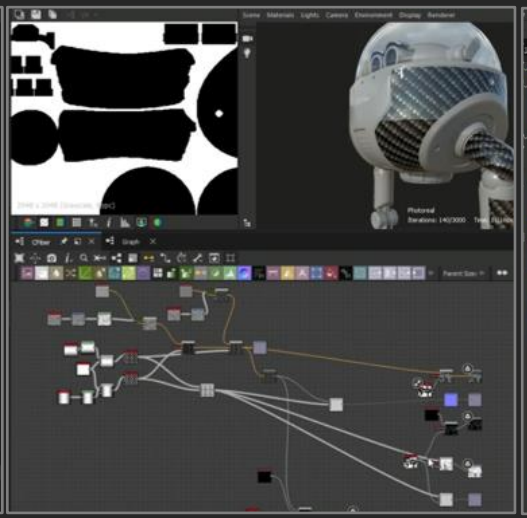

Author procedural materials, featuring RTX accelerated bakers

Substance Painter

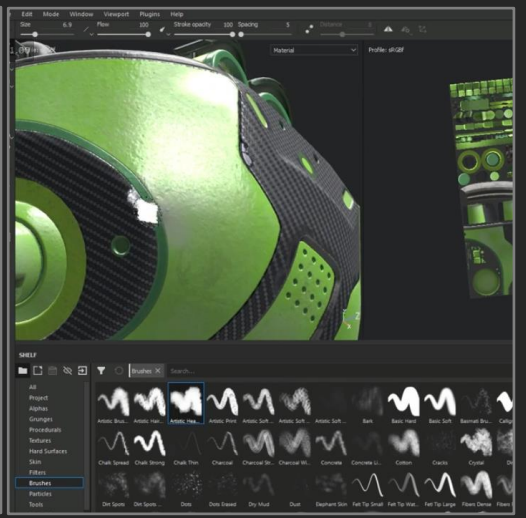

Paint materials onto 3D models, featuring RTX accelerated bakers

“ Our next generation of design tools, including Adobe Dimension, Substance Painter, Designer and Alchemist all benefit from hardware innovations like RTX that accelerate ray tracing and machine learning”

SEBASTIEN DEGUY

VP 3D & Immersive | Adobe

Adobe

CUDA-accelerated video & image processing across the suite

Premiere Pro

Edit and render video rapidly with GPU-accelerated effects

9.1x

Lightroom

Edit high resolution images smoothly with GPU-accelerated viewport

3.9x

Photoshop

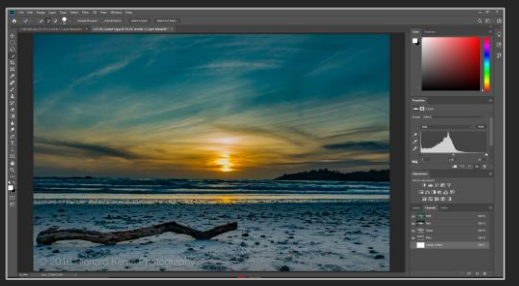

Edit quickly with 30+ GPU-accelerated features

1.3x

Illustrator

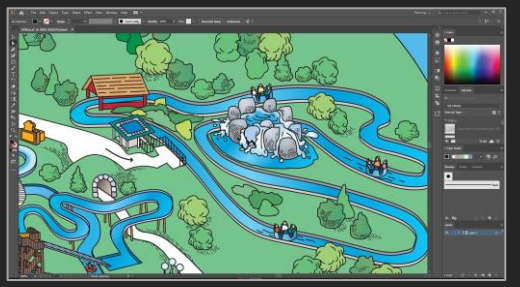

Pan and zoom smoothly with GPU-accelerated canvas

1.7x

Performance comparing RTX vs Core i9-9900K

Adobe Introducing RTX ray tracing in Adobe Dimension

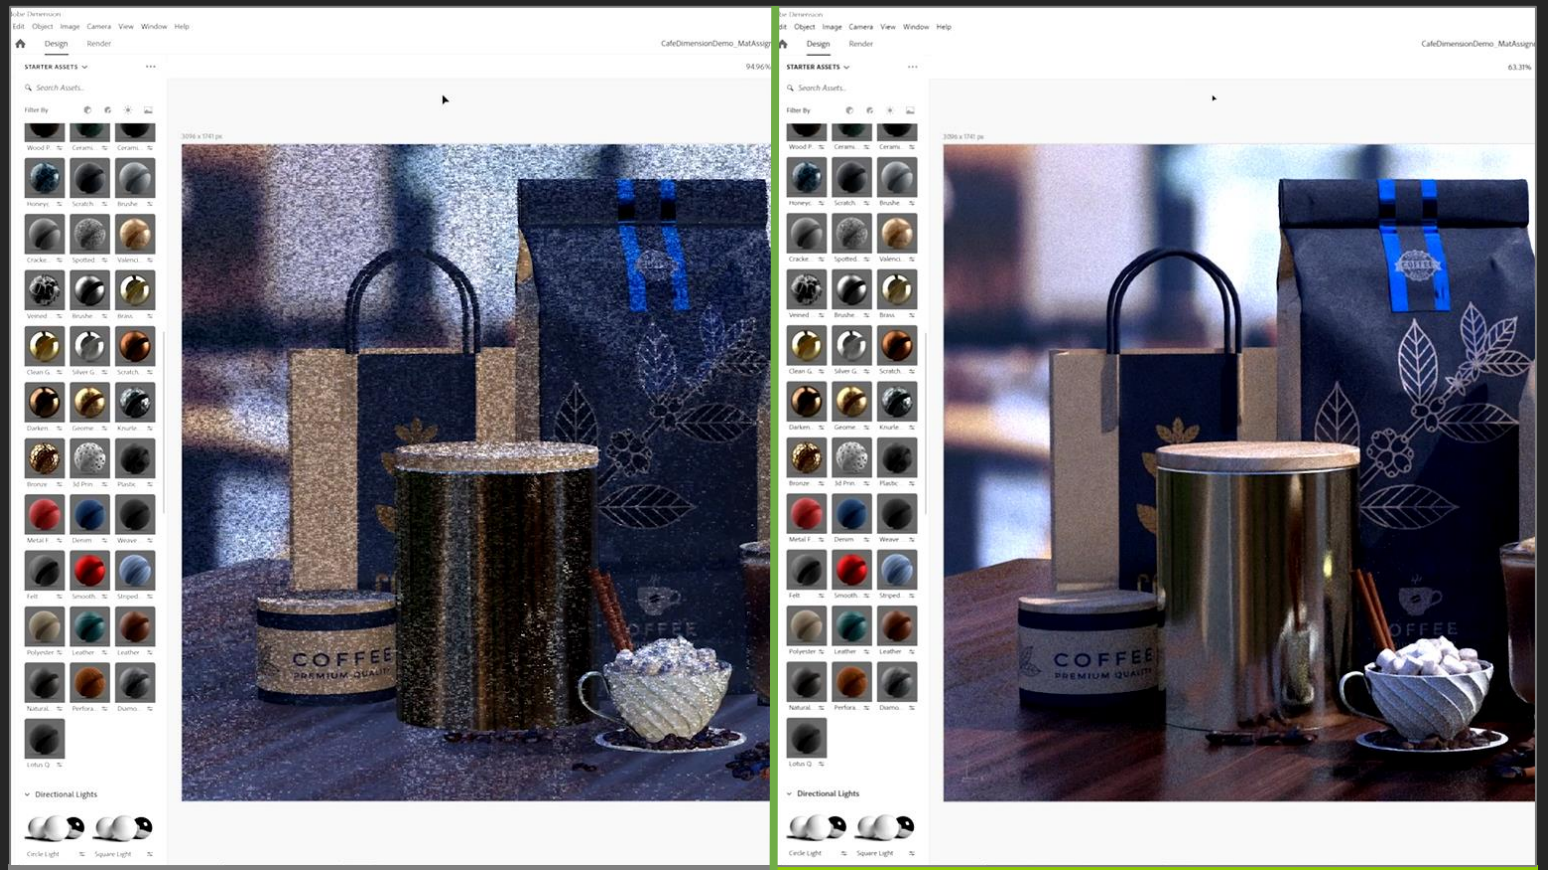

Performance comparing RTX vs Core i9-9900K

Adobe Introducing RTX AI Material Creation in Substance Alchemist

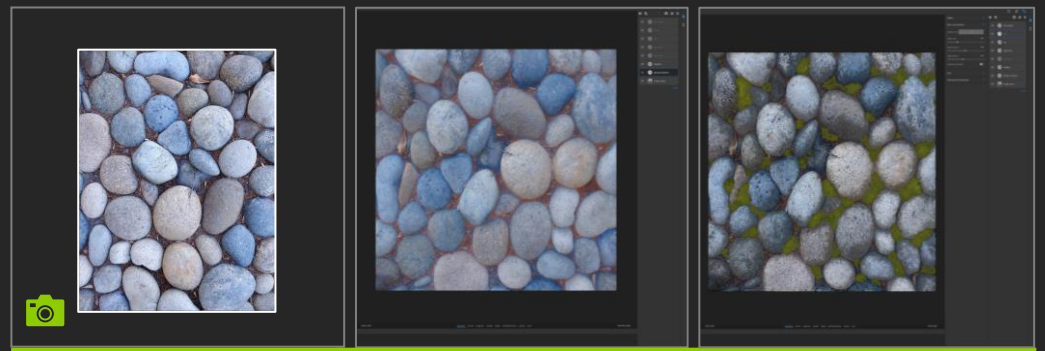

Take Photograph | Remove Shadows with AI | Add dirt & moss filters

RTX accelerated - available only on NVIDIA GPUs

*GPU-accelerated - up to **19X** faster than CPU*

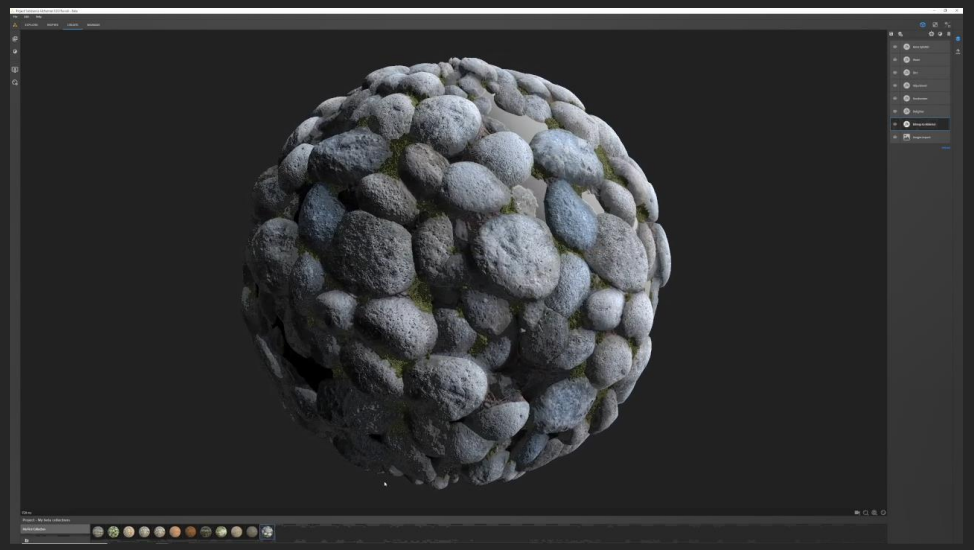

New Re-Lit 3D Material

Performance comparing RTX vs Core i9-9900K

Adobe Introducing GPU-Accelerated AI Auto Reframe in Premiere Pro

Traditional 16:9 landscape video

Portrait video

Portrait video with AI reframing - up to **4x** faster with a GPU

Performance comparing RTX vs Core i9-9900K

Adobe AI-Powered Neural Filters for Photoshop

- New AI workspace for artists to explore creative ideas & make amazing changes in seconds
- 7 of the 8 Neural Filters accelerated by NVIDIA RTX GPUs and Tensor Cores
- "Smart Portrait" Filter co-developed with Adobe using advanced NVIDIA StyleGAN2 technology

AUTODESK ARNOLD Version: 6

- Interactive GPU-accelerated ray tracing with RT Cores
- Enhanced image quality with Tensor cores for AI-accelerated denoising
- Ease of use with seamless switching from CPU to GPU rendering
- Render on multiple GPUs with NVIDIA NVLink™ for largest scenes
- Up to 17x faster than CPU rendering on Quadro RTX[1]

Rendering courtesy of Sohrab Esfehani

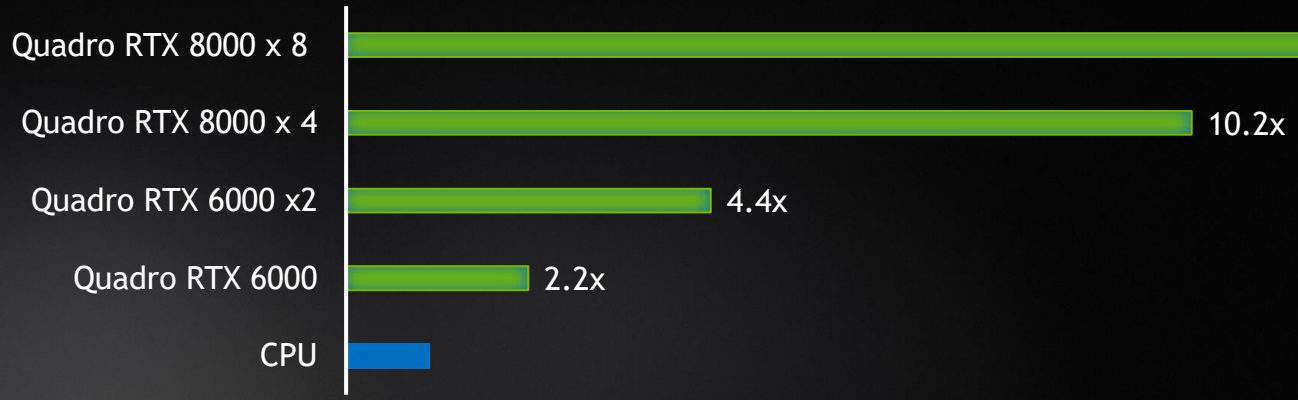

"We've worked closely with NVIDIA to optimize Arnold GPU to run on the latest RTX GPUs and RTX Server, and we're excited to get this latest update into the hands of new and existing Arnold customers."

Chris Vienneau
Senior Director, Media & Entertainment Products | Autodesk

Relative Performance. Tests run by NVIDIA on 2x Xeon Gold 6126 2.6 GHz/(6148 2.4 GHz for x8 results), 256 GB DDR4 RAM. Performance testing completed with Autodesk Arnold 6.0.1 Performance results may vary by scene.

Fast interactive rendering with Autodesk Arnold

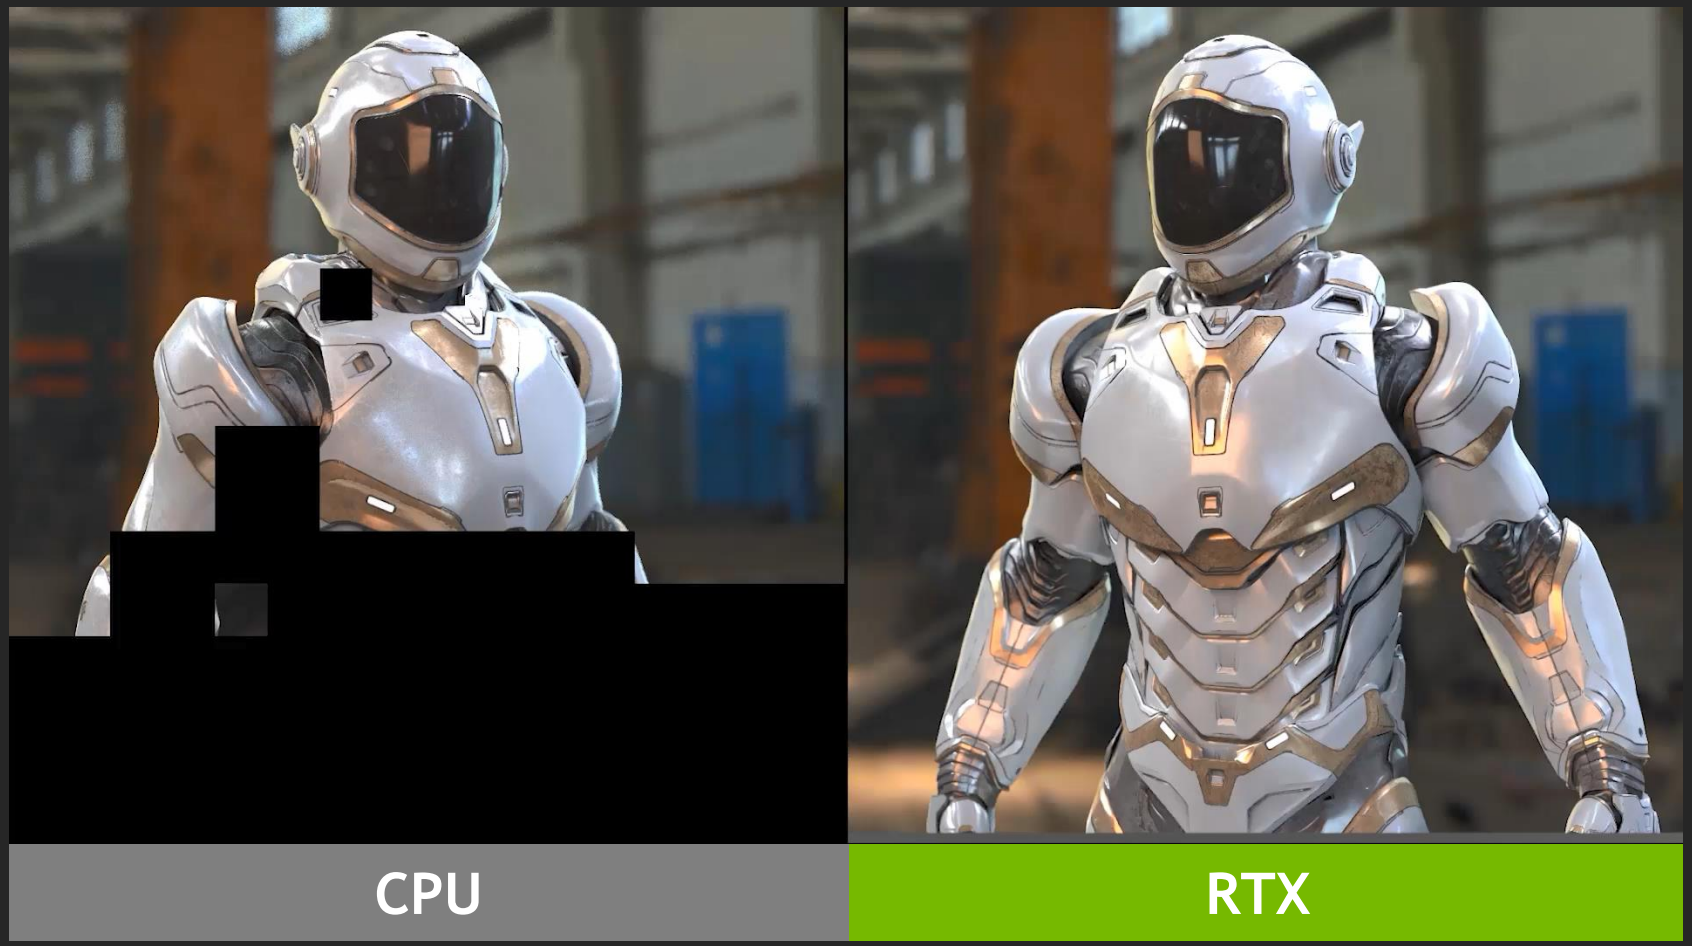

Performance comparing RTX vs Core i9-9900K

Version: 2.81

- Interactive GPU-accelerated ray tracing with RT Cores
- Ease of use switching from CPU to GPU rendering
- Scaled performance with multi-GPU configurations
- Up to 8x faster than CPU rendering on Quadro RTX[1]

Image courtesy of Blender

"In Blender Cycles we're always looking to reduce render time so artists can iterate faster. With NVIDIA RTX, core ray-tracing operations are now hardware accelerated by the GPU, making this the fastest version of Cycles yet."

Brecht Van Lommel
Lead Architect | Blender

Relative Performance. Tests run on 1x Xeon Gold 6154 3 GHz (3.7 GHz Turbo), 64 GB DDR4 RAM. Win10x64. Driver version 441.28. Performance testing completed with Blender 2.81. Comparison of aggregate total render time of 5 models. Performance results may vary depending on the scene.

V-Ray by Chaos Group

Version: V-Ray GPU Next

- Render high quality scenes faster with RTX accelerated ray tracing and AI-accelerated denoising
- 40% faster average GPU rendering on RTX vs. CUDA
- Render on multiple GPUs with NVIDIA NVLink™
- Up to 18X faster than CPU rendering on Quadro RTX[1]

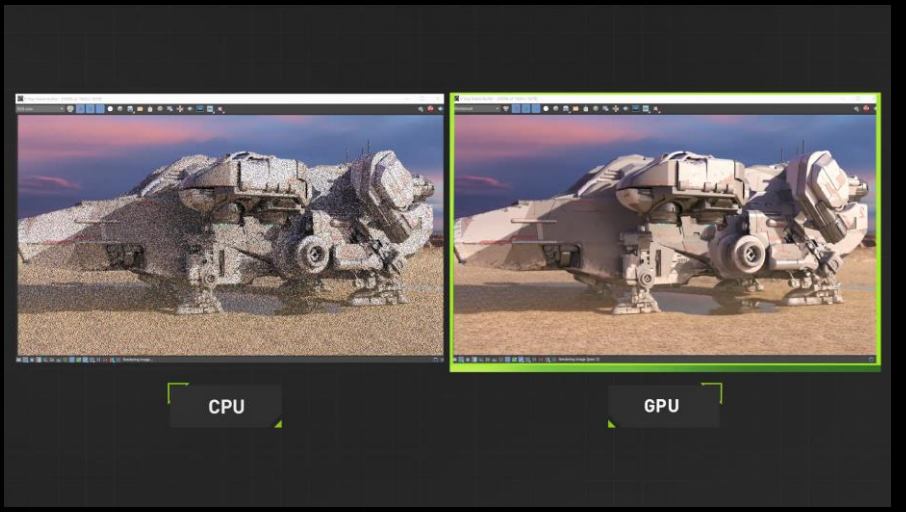

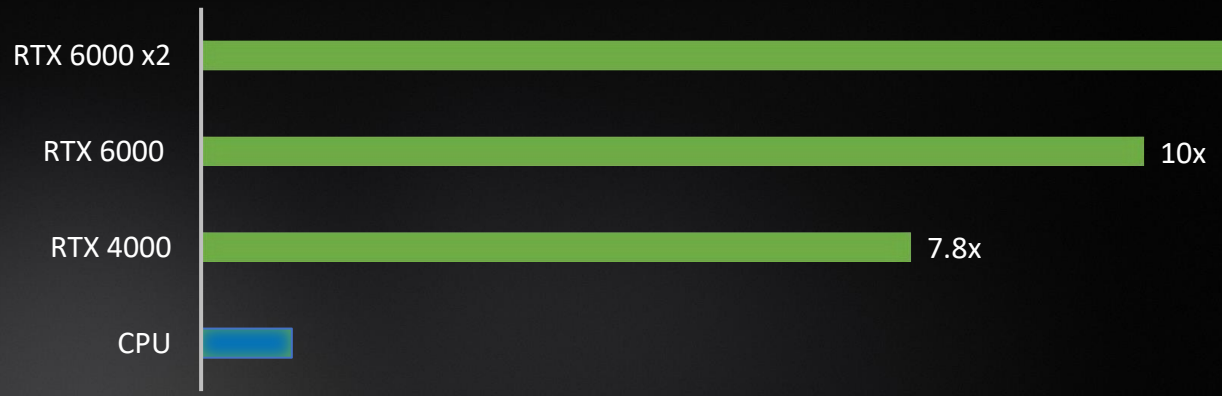

"Accelerating artist productivity is always our top priority, so we're quick to take advantage of the latest ray tracing hardware breakthroughs, by supporting NVIDIA RTX in V-Ray GPU, we're bringing our customers an exciting new boost in their GPU production rendering speeds."

Phillip Miller
Vice President of Product Management | Chaos Group

Relative Performance. Desktop Performance. Tests run on 1x Xeon Gold 6154 3 GHz (3.7 GHz Turbo), 64 GB DDR4 RAM. Win10x64. Driver version 441.28.

BLACKMAGIC RESOLVE Version: 16

- AI-accelerated video editing with Tensor Cores
- Increase footage res up to 4x with Super Scale
- Automatically color scenes, tag & track characters
- Learn styles from one clip and transfer to another
- Color grade and edit up to 14x faster on Quadro RTX[1]

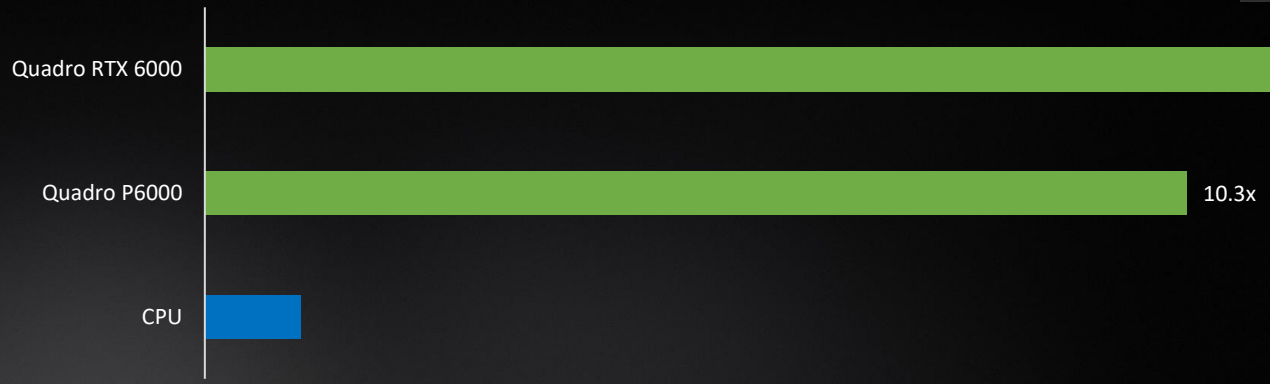

"DaVinci Resolve 16 Studio uses the latest multiple GPU innovations for AI and deep learning," "With the new DaVinci Neural Engine using NVIDIA CUDA 10 and TensorFlow acceleration, our tests show a better than 2x performance increase over previous NVIDIA GPUs. These same GPUs are also used for decoding and debayering Blackmagic RAW images which makes them an attractive investment.

Dan May
President| Black Magic Design USA

Relative performance. Performance testing conducted by NVIDIA in October 2019 on workstations equipped with 32GB RAM, Intel Core i9-9900K CPU with UHD360 graphics, Windows 10 x64, NVIDIA driver 431.86.
Davinci Resolve performance measures average FPS of video rendering using various effects such as face refinement, optical flow, blur, grain, noise reduction, at resolutions from 1080p to 8K. Performance results may vary depending on the scene.

AUTODESK® FLAME® 2021

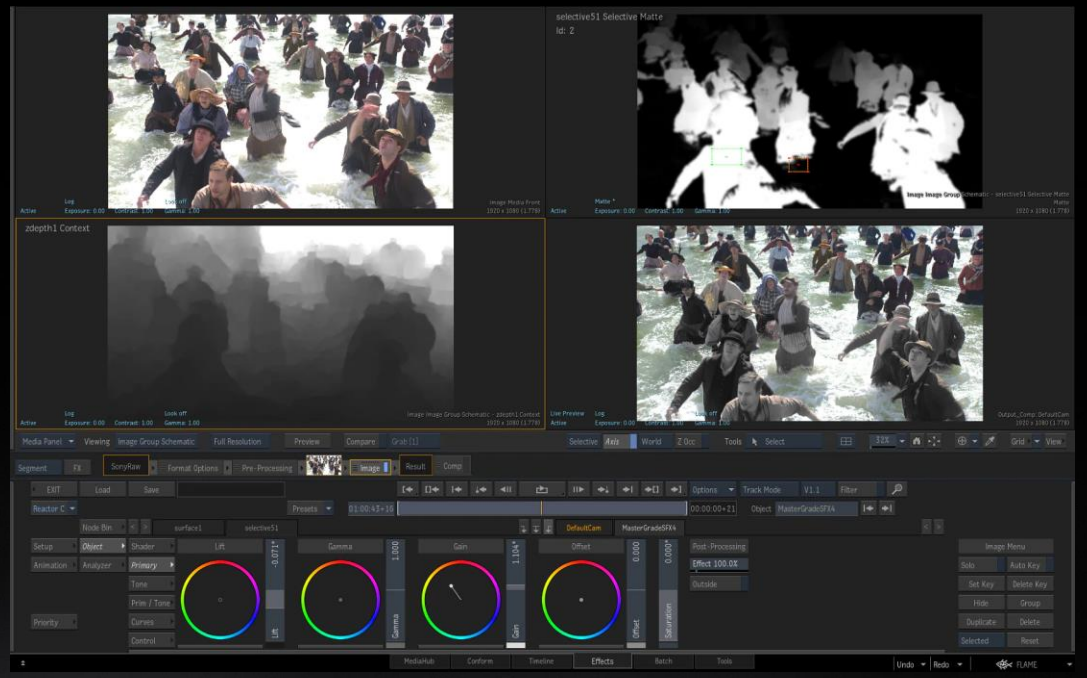

Z-depth map generator

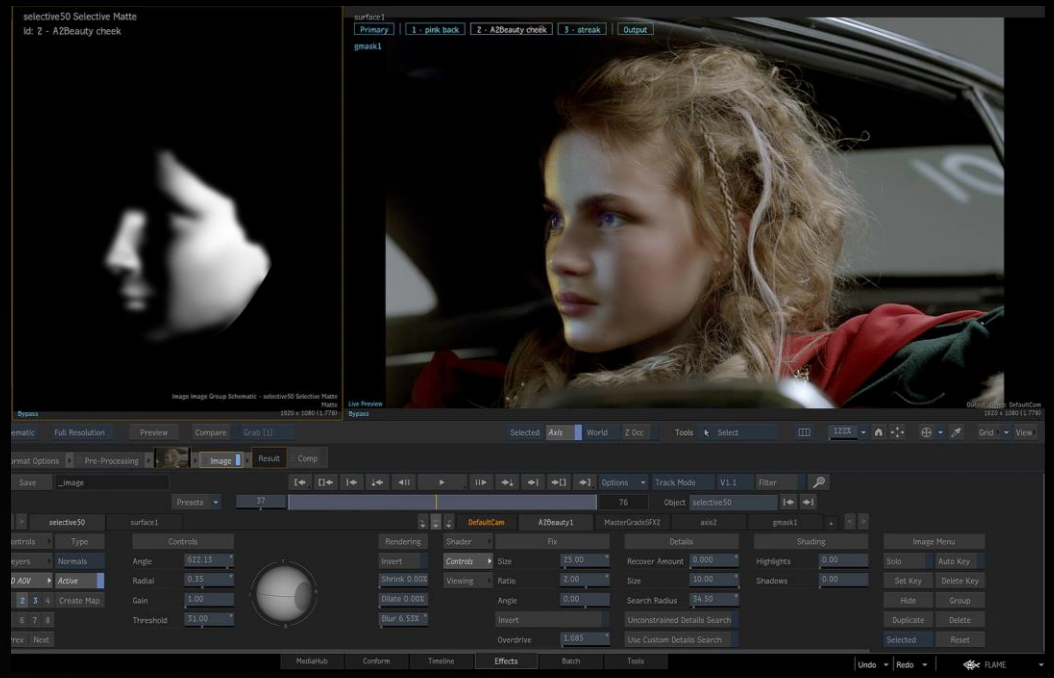

Human face normal map generator

THE EVOLUTION OF NVIDIA PROFESSIONAL VISUAL COMPUTING

Kepler

- 12 GB GDDR5 Memory
- 5.2 TFLOPs Single Precision performance
- 4K Display Resolution
- Multi-Display Technology with Quadro Sync

Maxwell

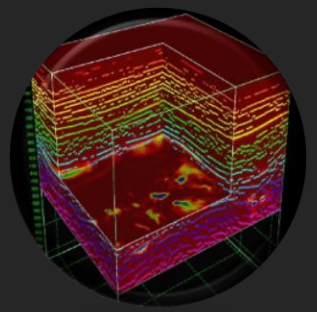

- 24 GB GDDR5 Memory
- 7 TFLOPs Single Precision performance
- Multi-Display Technology with Quadro Sync II
- Volume Rendering

Pascal

- 24 GB GDDR5X Memory
- 12 TFLOPs Single Precision performance
- 5K and 8K Display Resolution
- VR Ready

Turing

- RT Cores
- Tensor Cores
- 48 GB GDDR6 Memory with 672 GB/s Bandwidth
- 16.3 TFLOPs Single Precision performance
- Hardware-Accelerated Ray Tracing
- Hardware-Accelerated AI

Ampere

- 2nd Gen RT Cores
- 3rd Gen Tensor Cores
- 48 GB GDDR6 Memory with 768 GB/s Bandwidth
- 40 TFLOPs Single Precision performance
- Hardware-Accelerated Motion Blur
- Tensor Float 32 (TF32) Precision with Support for Structural Sparsity

2ND GEN NVIDIA RT CORES

Up to 2X the Ray Tracing Throughput of Turing

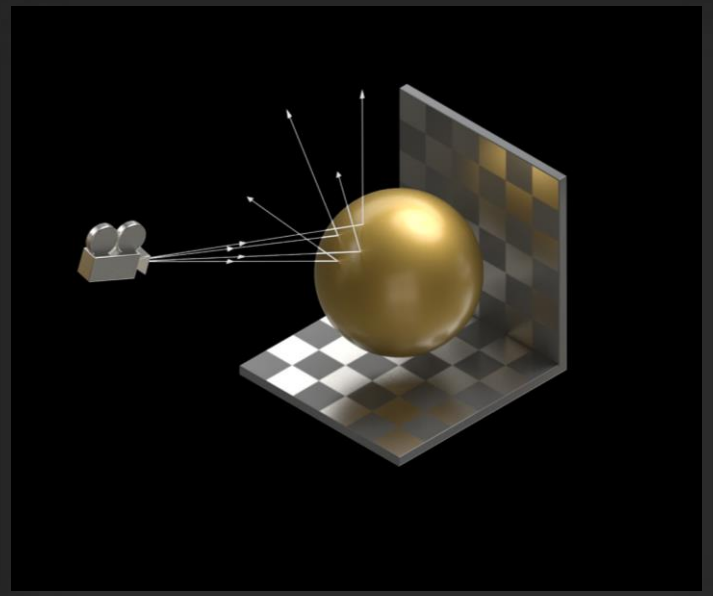

- Higher degree of interactivity with rendered scenes
- Accelerated motion blur rendering

Car Copter 25 model was provided by artist carbodyart and Cybertech Drone Pack was provided by artist IggyDesign through TurboSquid

*Performance measures gen to gen comparison of RTX 6000 to NVIDIA A40

3RD GEN NVIDIA TENSOR CORES IN NVIDIA A40

Faster model training, inferencing for AI & data science

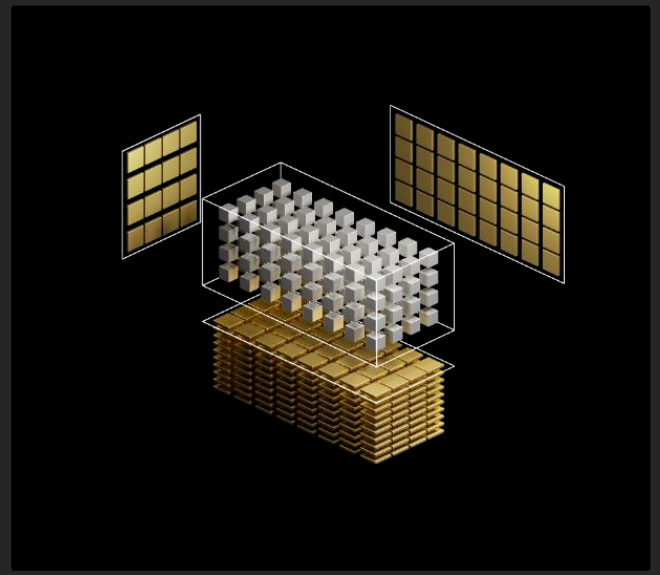

Up to 5X faster TF32 training

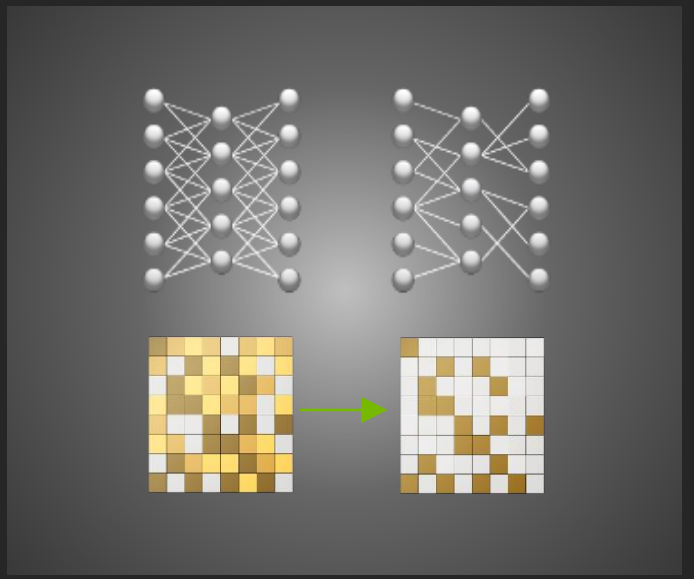

Up to 10X AI Performance with Sparsity

AI-Accelerated Graphics

*Performance measures gen to gen comparison of RTX 6000 to NVIDIA A40

PROFESSIONAL WORKLOADS

Up to 96GB of GPU Memory for Massive Models and Datasets

Movie-quality virtual production, scalable visualizations with 3 available display ports

Life-like renders and product designs

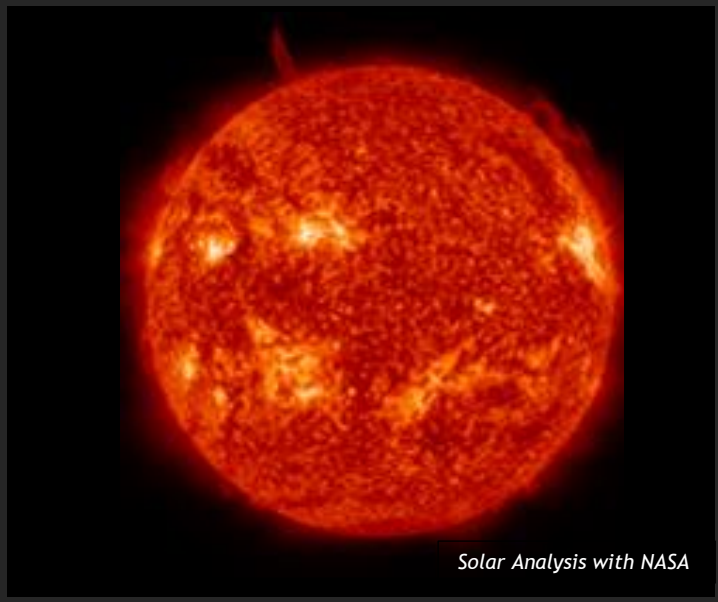

Actionable data science from big data

*NVIDIA A6000 and A40 features 48GB of GPU Memory, scalable up to 96GB with 2-way NVLink

RTX A6000

NVIDIA RTX A6000

Workstation Performance Amplified

NVIDIA Ampere CUDA Cores
Up to 2X FP32 throughput of previous generation

2nd Generation RT Cores
Up to 2X throughput of previous generation

3rd Generation Tensor Cores
Up to 5X* (10X with sparsity) throughput with TF32

48 GB GDDR6 Memory
Largest frame buffer for professional graphics card

PCIe Gen 4
2X bandwidth of PCIe Gen 3

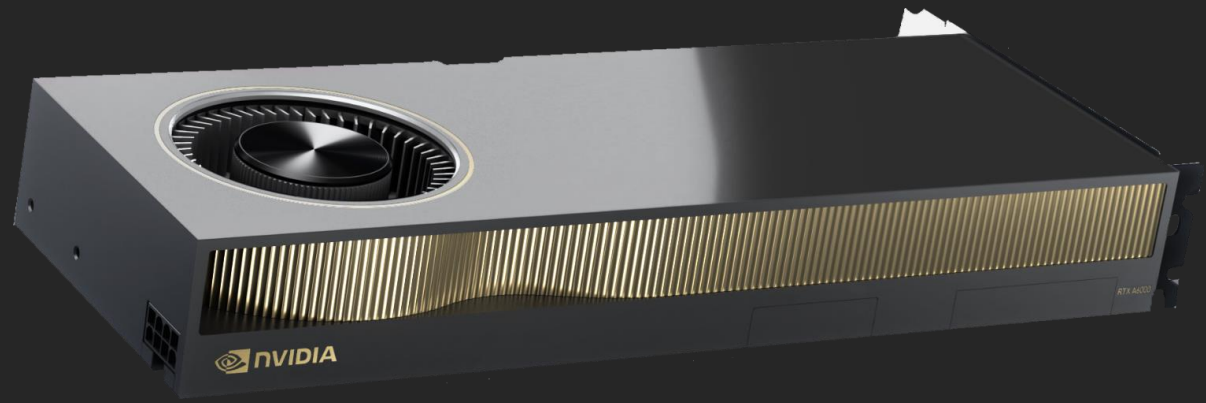

*Compared to the previous generation RTX6000

NVIDIA A40

NVIDIA A40

World's Most Powerful Data Center GPU for Visual Computing

NVIDIA Ampere Architecture CUDA Cores
Up to 2X FP32 throughput of previous generation

2nd Generation RT Cores
Up to 2X throughput of previous generation

3rd Generation Tensor Cores
Up to 5X throughput with TF32 (10X with Sparsity)

48 GB GDDR6 Memory
Largest frame buffer for professional graphics

PCIe Gen 4
2X bandwidth of PCIe Gen 3

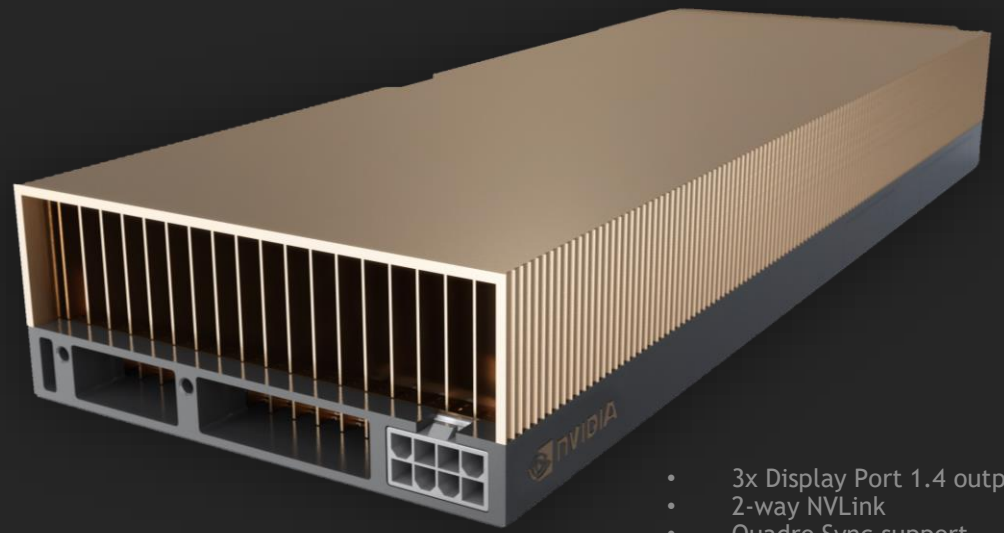

- 3x Display Port 1.4 outputs
- 2-way NVLink
- Quadro Sync support
- vGPU software support (early 2021)
- Hardware secure boot

*Performance measures gen to gen comparison of RTX 6000 to NVIDIA A40

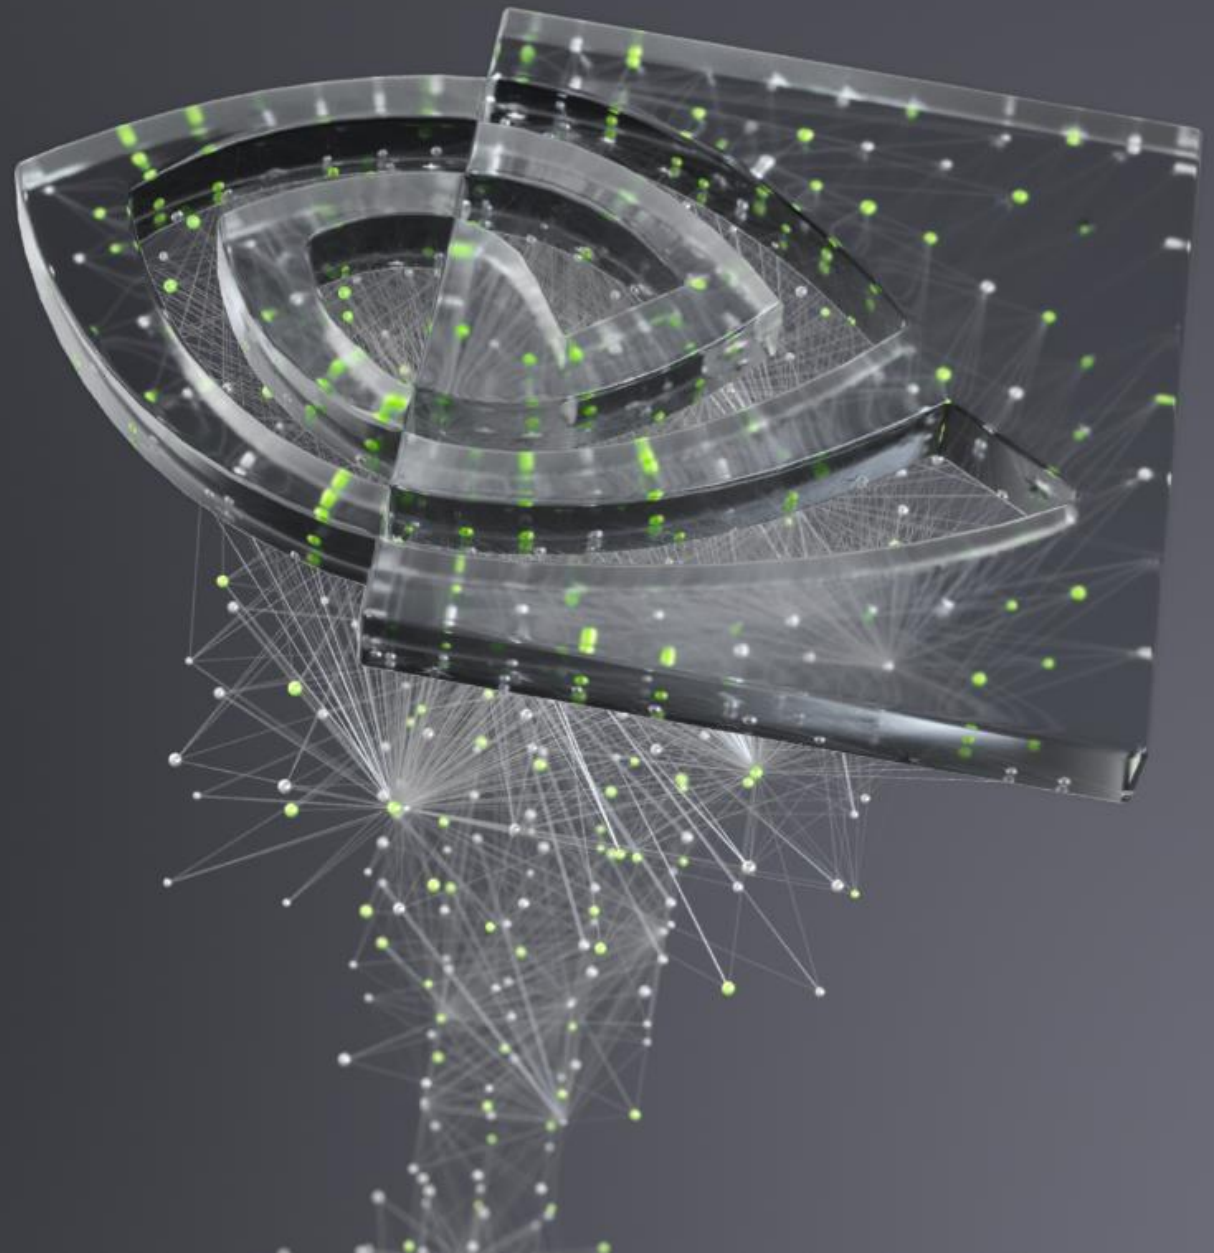

ENTERPRISE MEDIA SOLUTIONS

NVIDIA

OMNIVERSE

Overview & Technology

NVIDIA

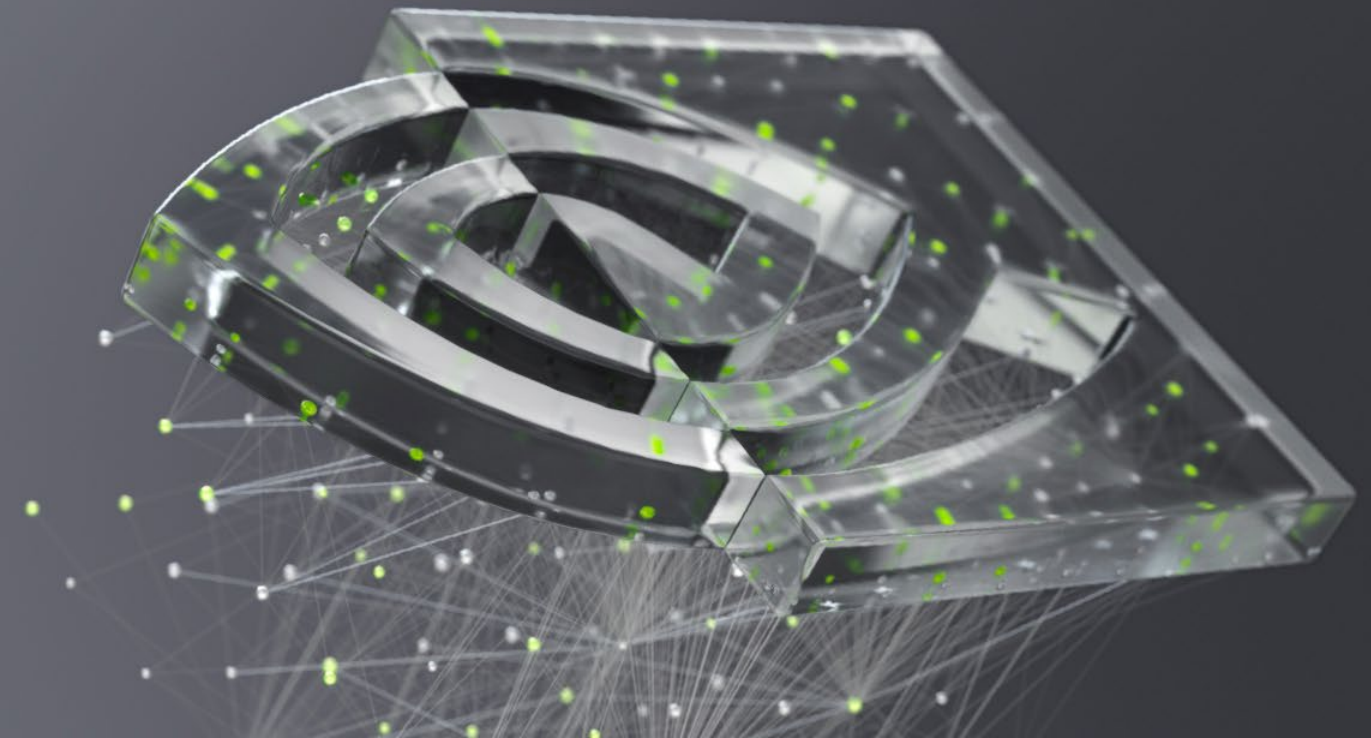

NVIDIA TECHNOLOGIES

UNREAL ENGINE

AUTODESK REVIT

SUBSTANCE by Adobe

unity

AUTODESK® MAYA®

TOOLS

NVIDIA TECHNOLOGIES

NVIDIA OMNIVERSE

UNREAL ENGINE

AUTODESK REVIT

SUBSTANCE by Adobe

unity

AUTODESK® MAYA®

TOOLS

SIMULATING REALITY

AEC

Design

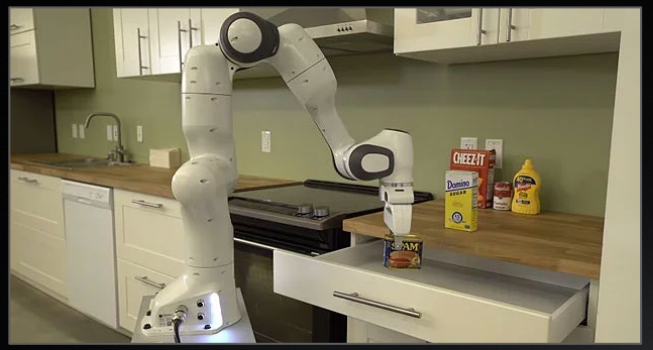

Robotics

Autonomous Vehicles

Media & Entertainment

NEW CHALLENGES

Content Pipeline

NEW CHALLENGES

Realtime Ray Tracing

NVIDIA OMNIVERSE™

Top Industry Tools

+

Realtime Collaboration

+

Open Standards

CONNECTING ALL THE TOOLS TOGETHER THROUGH USD

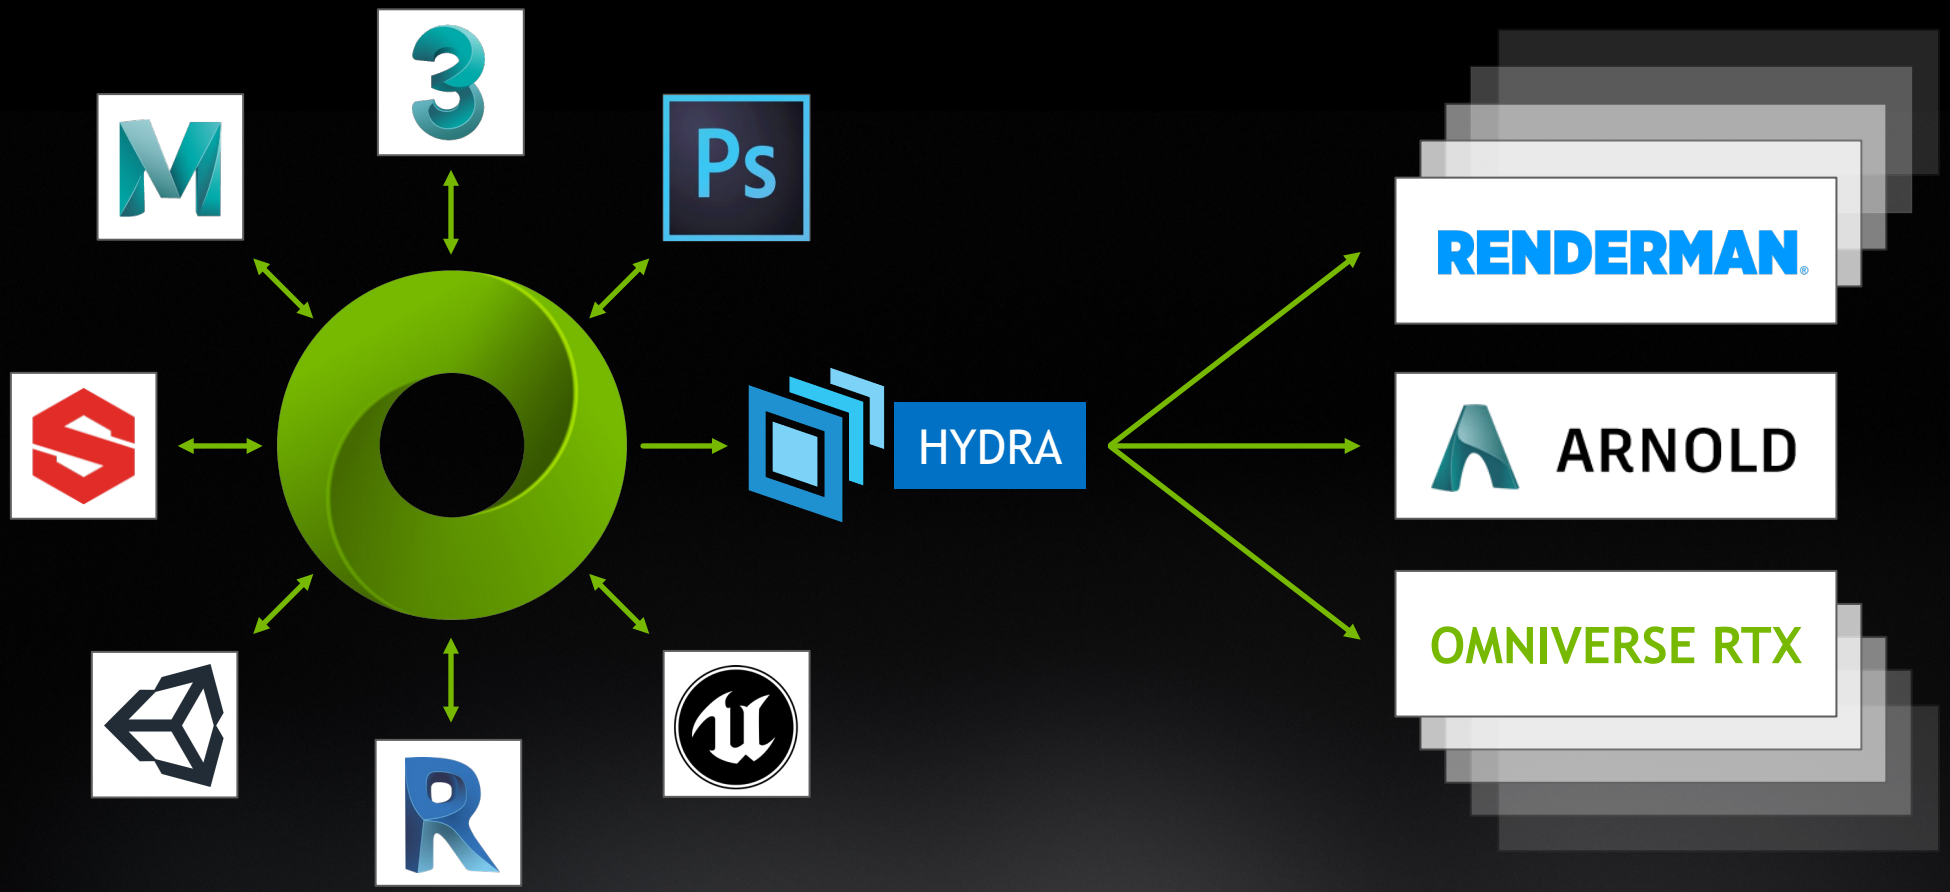

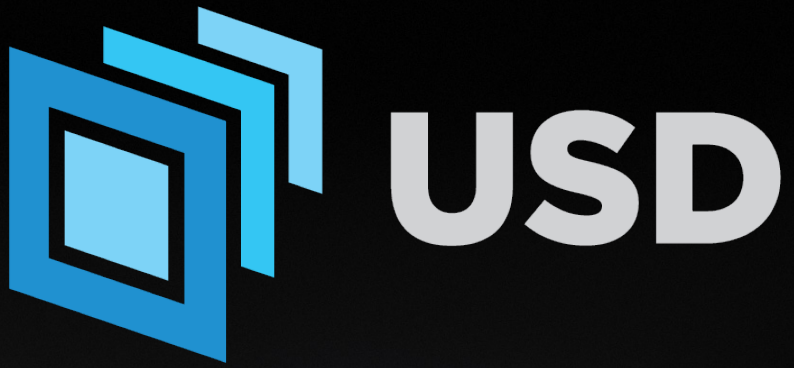

UNIVERSAL SCENE DESCRIPTION

Open sourced API and format for complex scene graphs

Built for interchange between apps with wildly varying scene representations

Introduces novel concept of layering

Allows for large teams with different department working on the same scene simultaneously

Quickly becoming a standard—even outside of M&E

We believe USD will be the HTML of 3D virtual worlds

OMNIVERSE

CONNECT	NUCLEUS	CREATE	SIMULATION	RTX RENDERER
		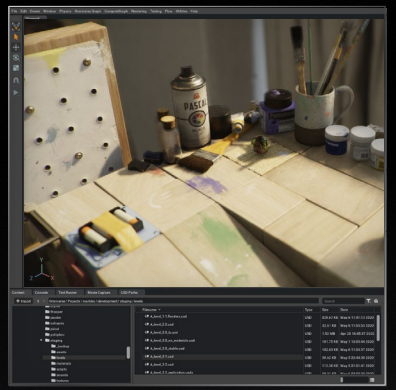		
Connection SDK / Plugins	Core Services / On Prem / Cloud	Viewer / Editor / Framework	Physics / AI / Animation / Behavior	Realtime / Scalable / Accurate / MDL

NUCLEUS

Essential services for collaboration

Asset Database

USD native

Versioning

Overlays on existing storage infrastructure

Incremental Updates (live sync)

Asset tagging (manual and AI-based classification)

Search

Microservice infrastructure (OpenFAAS, K8s, …)

File format conversion

LoD generation

Thumbnail/Batch rendering

Web interface (like Gdrive, Dropbox, …)

STANDALONE

Artist Workstation

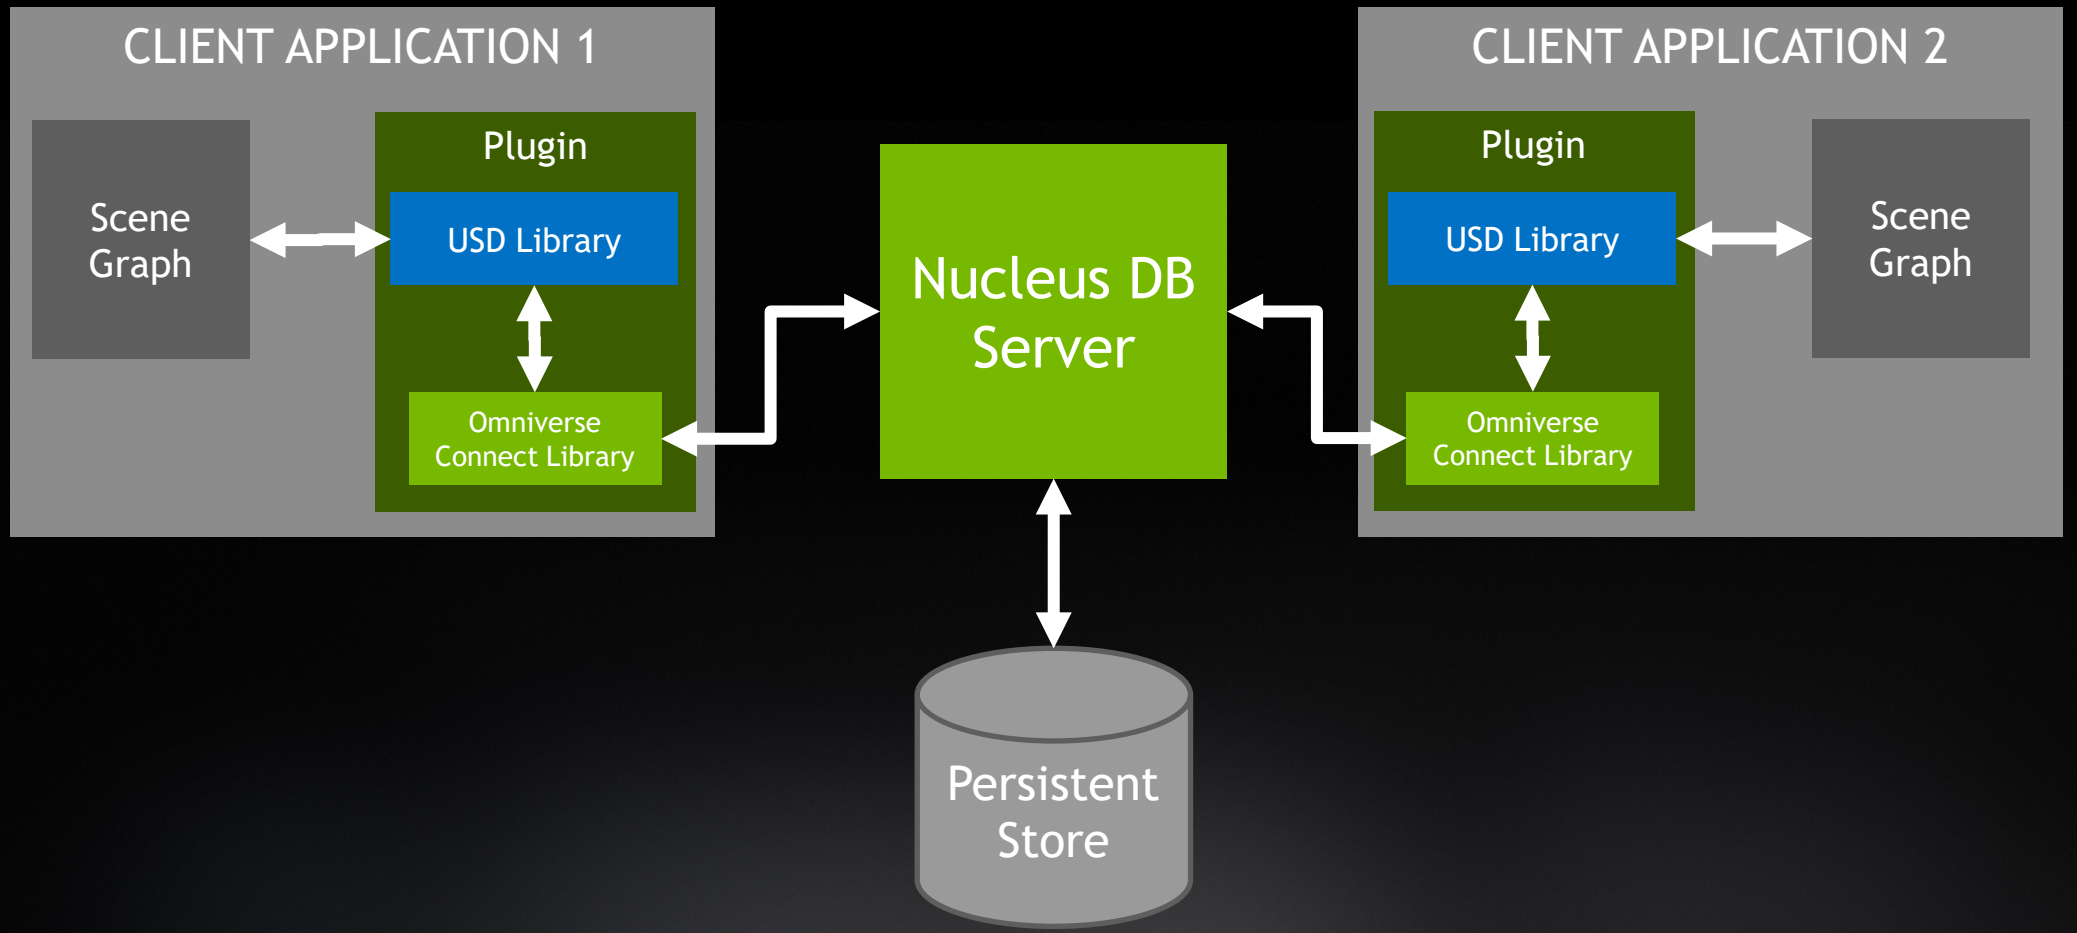
CLIENT APPLICATION 1
Scene Graph
Plugin
USD Library
Omniverse Connect Library
Nucleus DB Server
Persistent Store
CLIENT APPLICATION 2
Plugin
USD Library
Omniverse Connect Library
Scene Graph

ACROSS TEAMS

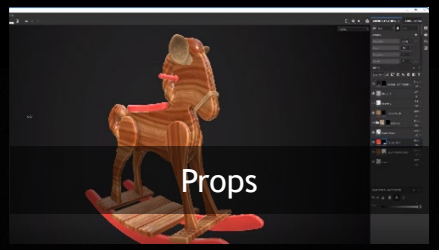

Los Angeles

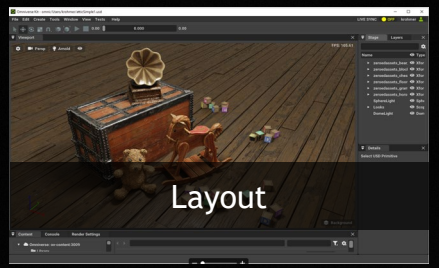

London

Bangalore

ANYWHERE ON ANY DEVICE

Stream Realtime Ray Tracing

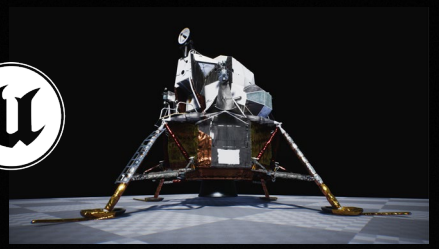

OMNIVERSE
RTX

Streaming Rendering

Virtual Camera
using ARKIT

MDL IN USD

Official Schema Based on USDShade | USDView Plugin for Previewing
Open Source MDL SDK for Renderer/App Integration

```
def Material "flex_material" {
...
def Shader "flex_material" {
      uniform token info:implementationSource = "sourceAsset"
      uniform asset info:mdl:sourceAsset = @nvidia/core_definitions.mdl@
      uniform token info:mdl:sourceAsset:subIdentifier =
      "::nvidia::core_definitions::flex_material"
```

REALTIME TO OFFLINE

Build Once

USD

MDL

Render Anywhere

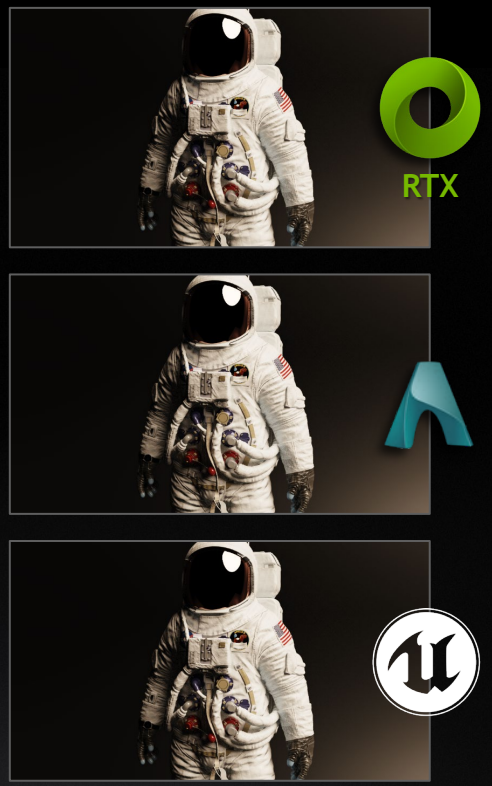

SCALABLE REALTIME RENDERING

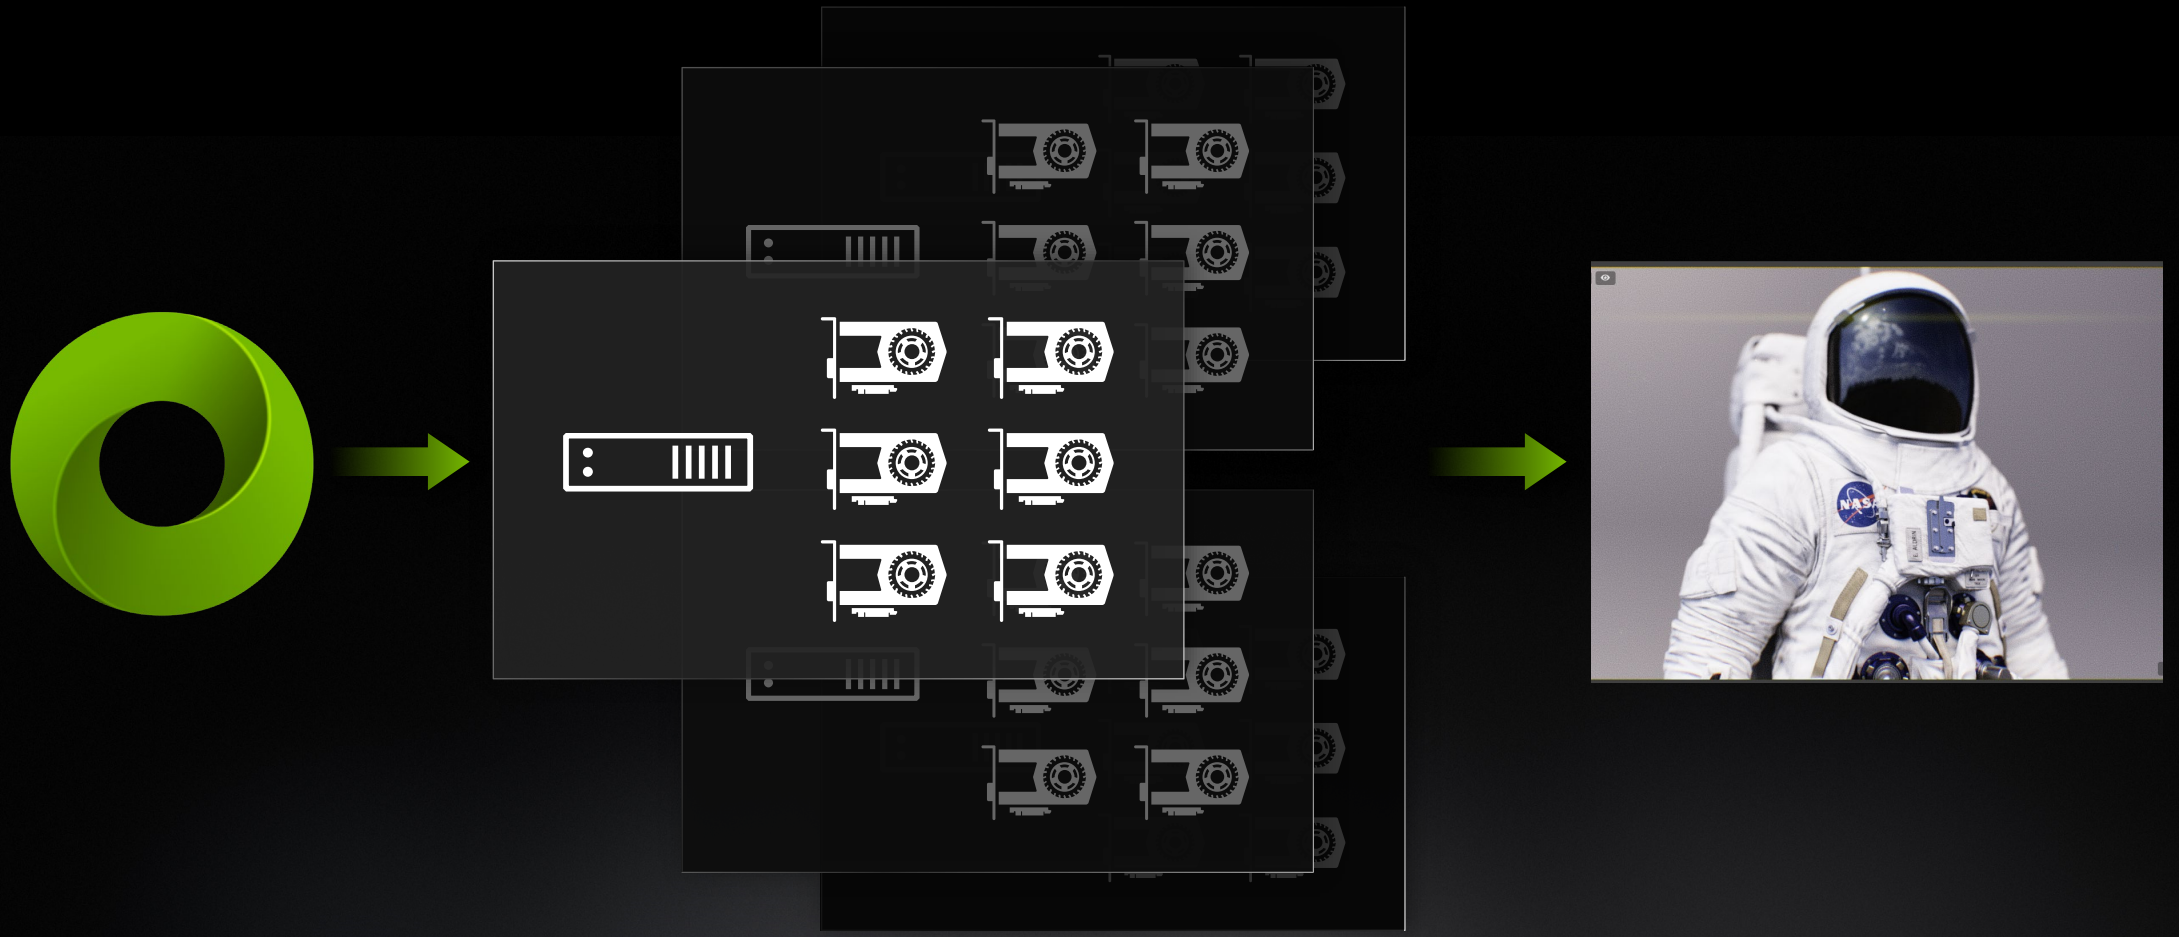

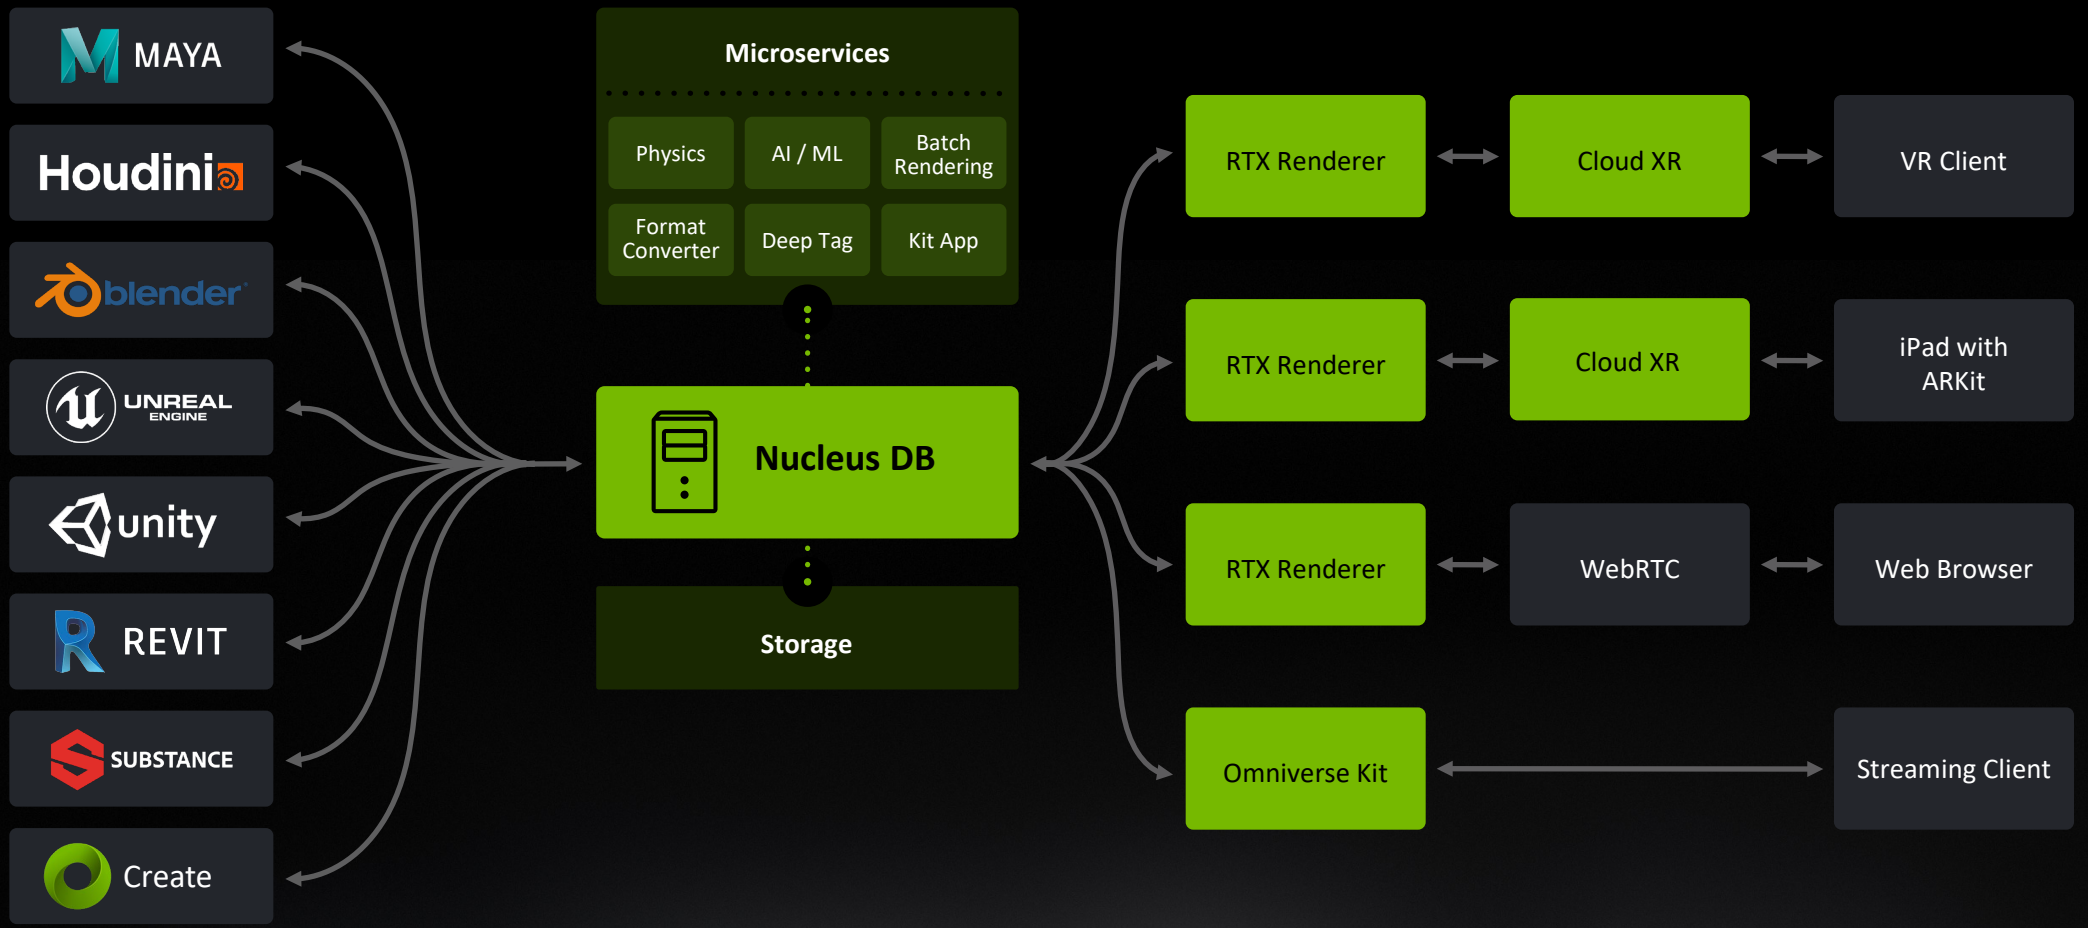
MAYA
Houdini
blender
UNREAL ENGINE
unity
REVIT
SUBSTANCE
Create
Microservices
Physics
AI / ML
Batch Rendering
Format Converter
Deep Tag
Kit App
Nucleus DB
Storage
RTX Renderer
Cloud XR
VR Client
RTX Renderer
Cloud XR
iPad with ARKit
RTX Renderer
WebRTC
Web Browser
Omniverse Kit
Streaming Client

KIT STACK

Developer View

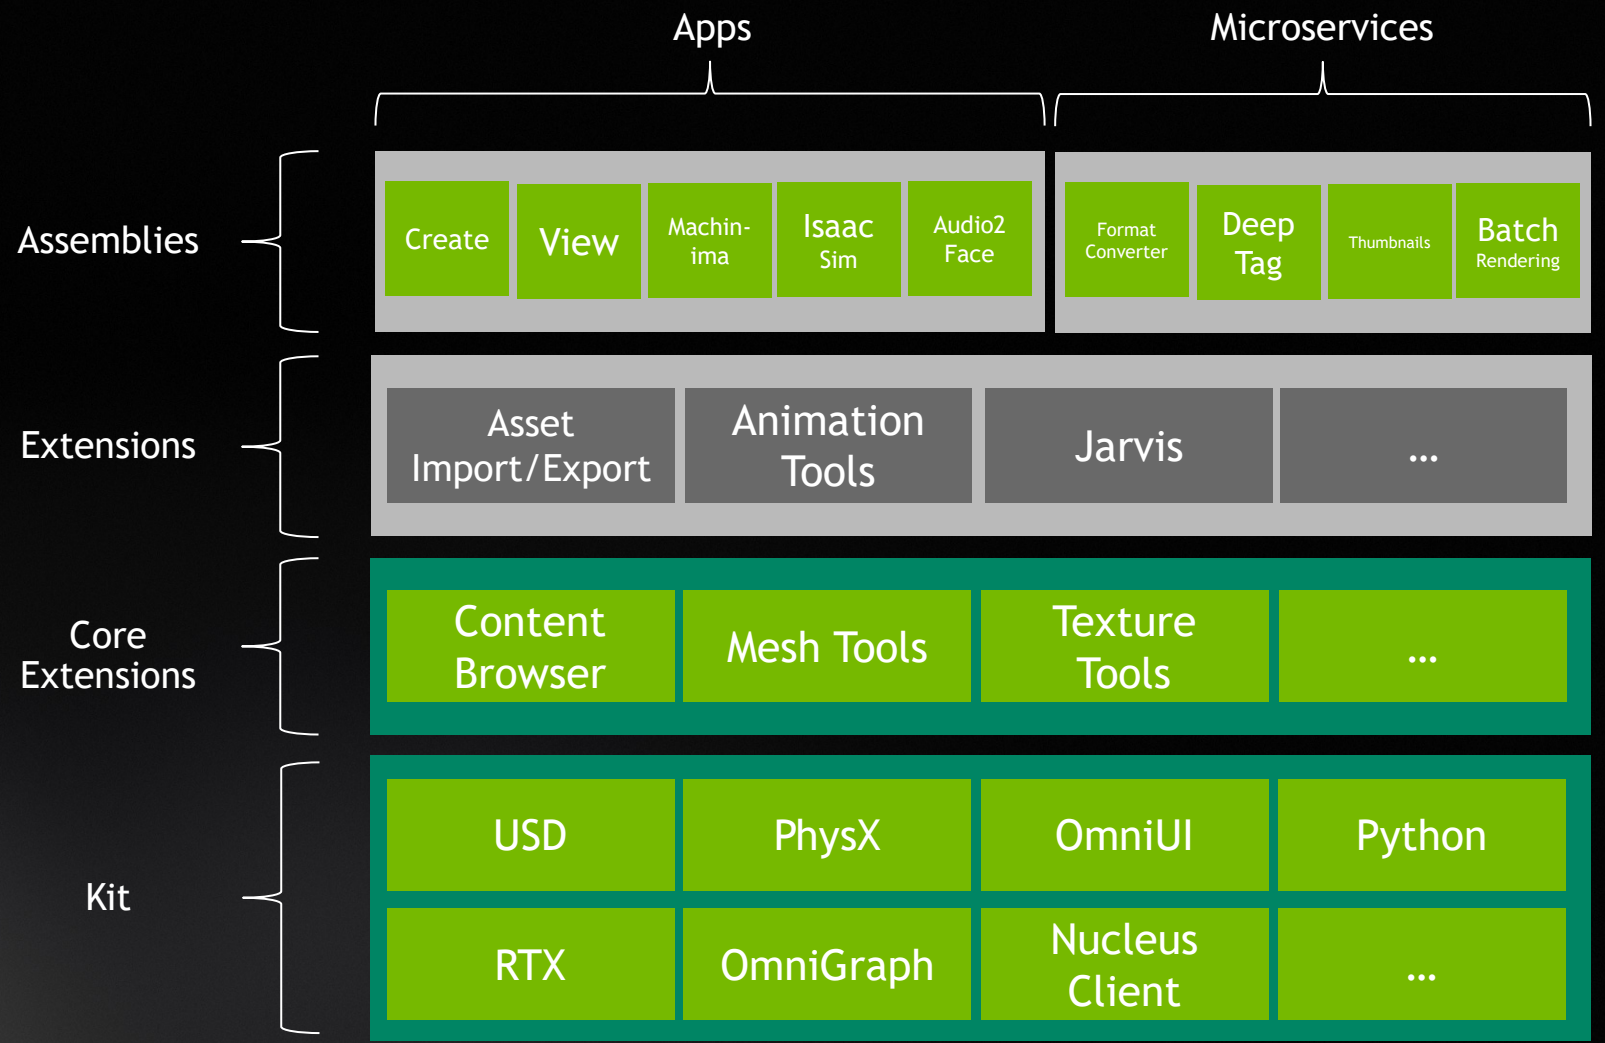

WHAT'S NEXT

Currently in early access: 2020.2

2020.3 coming soon

Sign up Here

https://developer.nvidia.com/nvidia-omniverse

Learn about USD at

http://usd.nvidia.com

NVIDIA

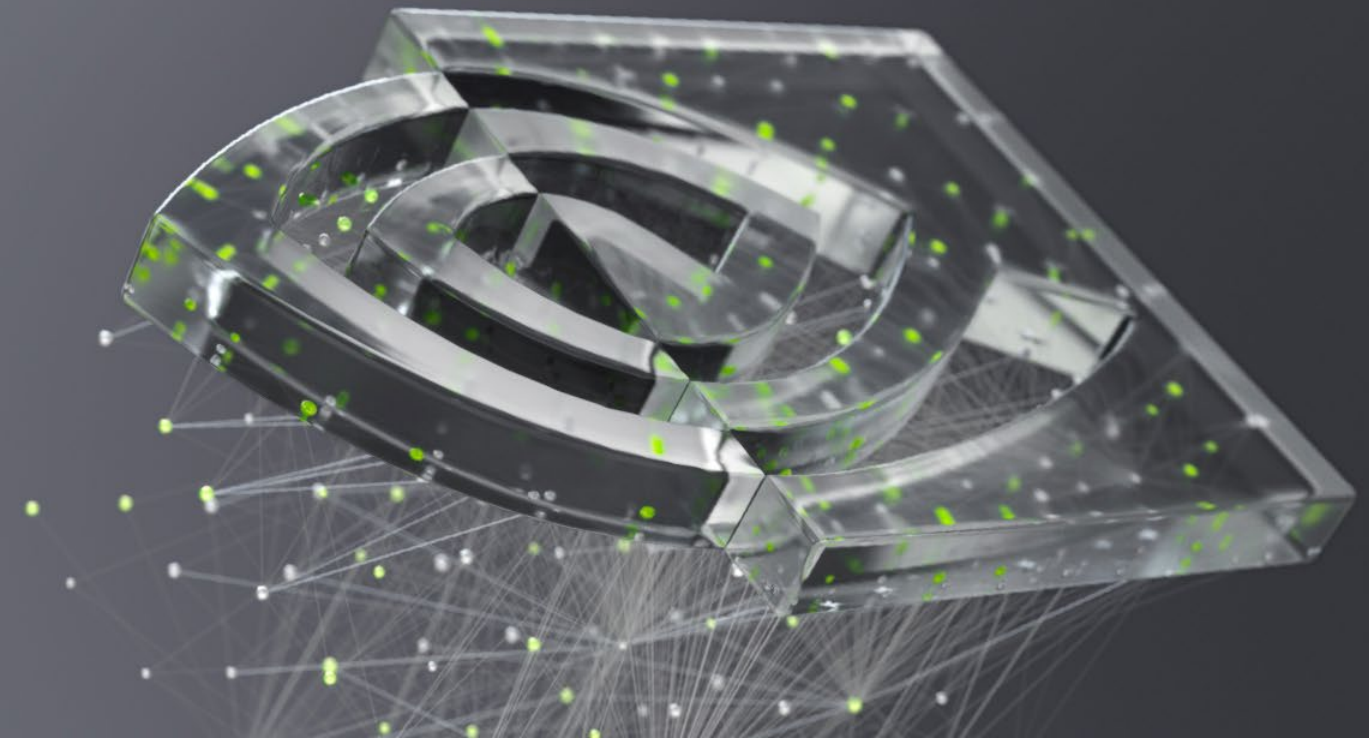

Thank you for attending!

FOR MORE INFORMATION CONTACT:

gopny@pny.com

NVIDIA QUADRO
AUTHORIZED PARTNER
PNY