



# NVIDIA Quadro RTX

The Fusion of Graphics and AI

S9969 | Booth 1133




**RTX**



EXPLORE **MORE.**  
DELIVER **MORE.**


# NVIDIA Quadro RTX

The fusion of graphics and AI




## RT Cores deliver real-time ray tracing

- Model how light and materials behave in 3D
- 25x faster than Pascal GPUs
- Renders film effects at 30x CPU speeds



## Advanced shaders for graphics and VR

- Mesh shading
- Variable rate shading
- Texture space shading



## Tensor Cores power AI-augmented applications

- AI denoising
- Resolution scaling
- Video re-timing



# NVIDIA QUADRO



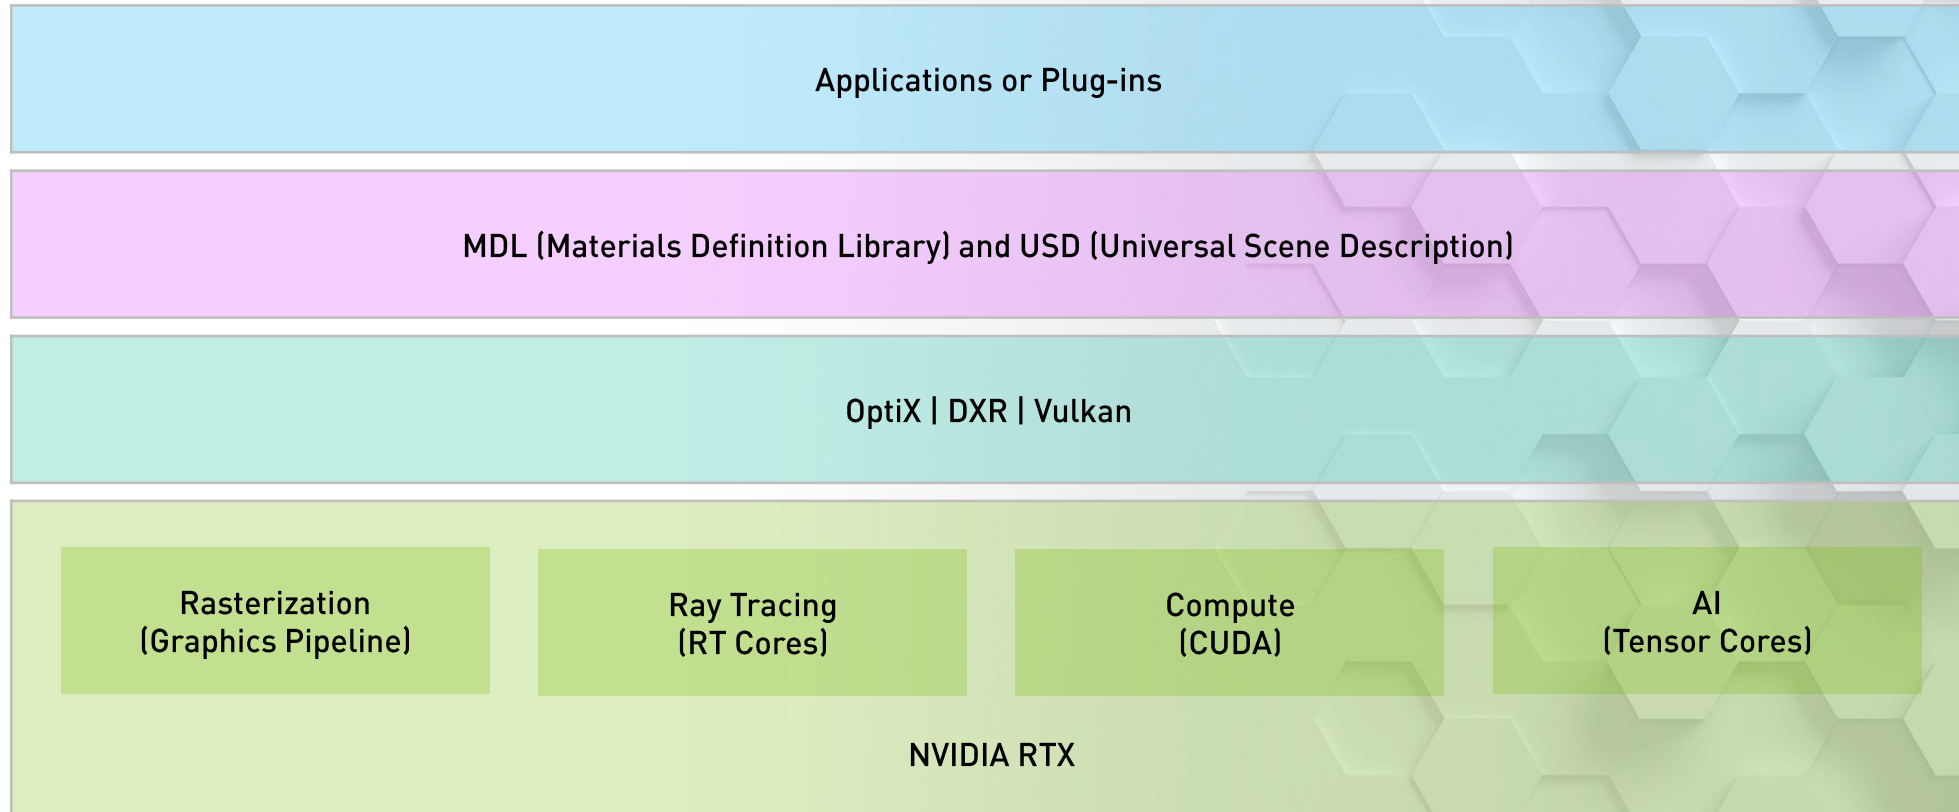
## RTX

**PNY**  
PNY PRO

EXPLORE **MORE.**  
DELIVER **MORE.**

# NVIDIA RTX Technology

Next generation hybrid rendering




# NVIDIA RTX RT Cores

Deliver hardware accelerated ray tracing

## Functionality includes:

- Ray-triangle intersection checks
- Bounding Volume Hierarchy (BVH) manipulation
- Real-time performance in application viewports



American Gods image courtesy of Tendril



RTX

PNY  
PNY PRO

EXPLORE MORE.  
DELIVER MORE.

# NVIDIA RTX Advanced Shaders

More objects per scene, flexible control over detail and performance, better VR

## Mesh Shading




- Eliminates CPU call bottlenecks and draws triangles more efficiently

## Texture Space Shading




- Decouples shading from screen space, improving shading efficiency and reuse

## Variable Rate Shading



- Pixel shading rate control for effects like motion, blur, foveated rendering



## Multi-View Rendering



- Extends Single Pass Stereo with unique view origin positions or directions

# NVIDIA RTX Tensor Cores

Next generation hardware accelerated deep learning




## Key benefits include:

- Hardware acceleration of deep learning enabled tools and applications
- Support for additional precision modes for improved performance
- Turing Tensor Cores tuned for fast training and inferencing performance

# NVIDIA Quadro RTX VR

New capabilities for ultimate VR experiences



## Variable Rate Shading

- Match lens optics
- Foveated rendering
- Places detail where most needed
- Context sensitive

## Single Pass Stereo 2.0

- Even more efficient GPU use
- SMP for 4 independent views
- Wide FOV HMD's
- Wrap-around multi displays

## VirtualLink

- Open industry standard
- Single cable power, data, video
- Utilizes USB-C Alt-Mode
- Easy VR setup, reduced cabling

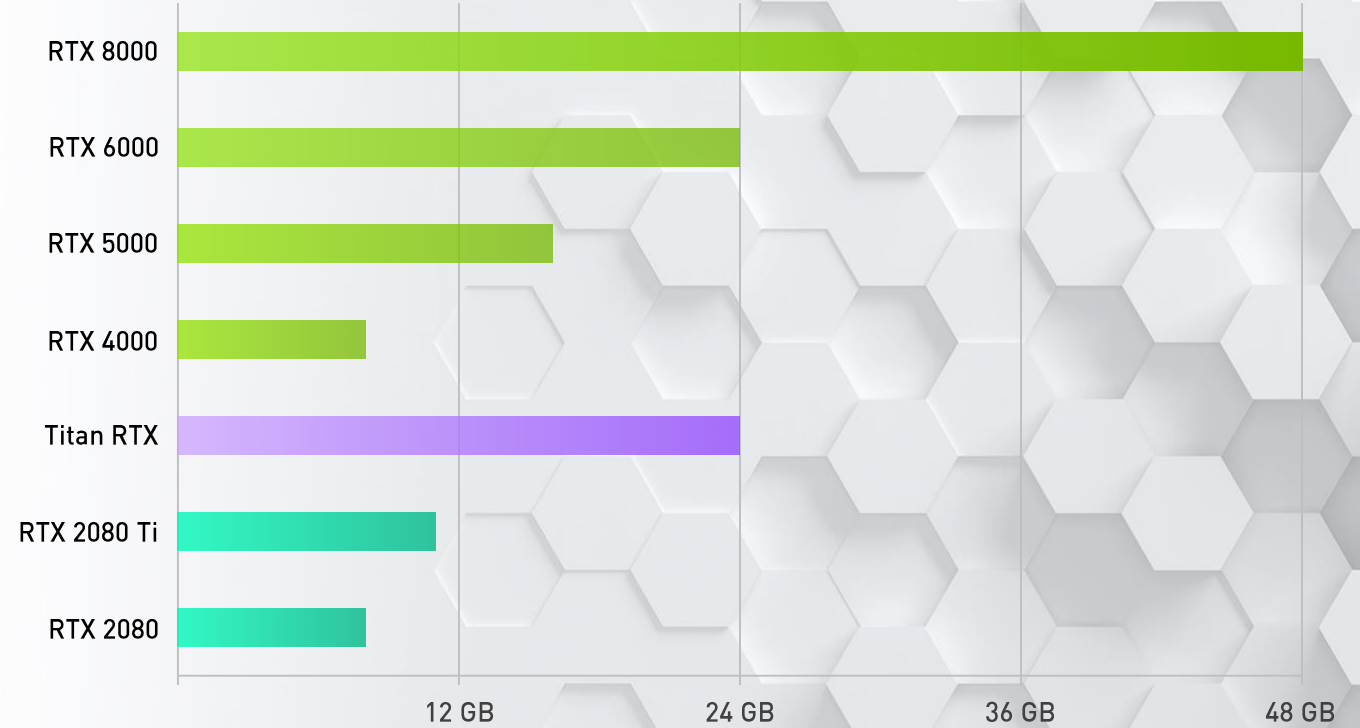
## Performance

- Rasterization
- Ray tracing
- Physically-based audio
- Dual-input VR-SLI

# NVIDIA Quadro RTX

Provides the memory required for professional VR

Geometry	See entire urban scenes, facilities, airplanes and cars in VR with full fidelity
Textures	Use more and larger textures to create the most realistic VR environments possible
Panoramas	Move around without interrupting the VR experience by using multiple panoramas
Light Fields	Provides most realistic VR experience possible using photorealistic imagery







Gamers can cut down on fidelity and still play. Professional customers need **full data fidelity** to make the right decisions and create the best designs.



# NVIDIA Quadro RTX Value Proposition

Spans all key markets and solutions

	M&E	Manufacturing	AEC	
Rendering	 Content Creation	 Product Design	 Building Design	Realtime rendering speeds up creative workflows
AI	 In-Painting, Upres	 Generative Design	 Generative Design	AI-augmented tools accelerate the creative process
VR	 Content Creation	 Design Review	 Design Review	VR content creation and design reviews

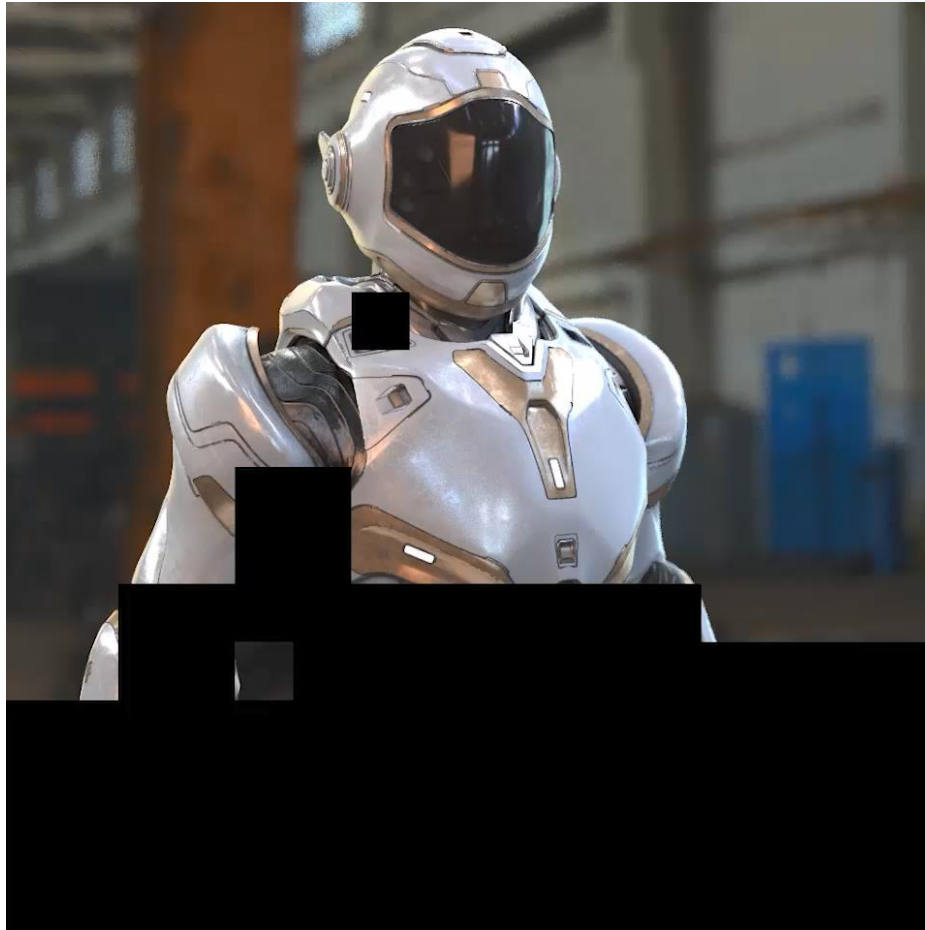


RTX



RT

# NVIDIA Quadro RTX | Create More, Wait Less



CPU



RTX

Note: CPU Core i9-7900X, GPU NVIDIA RTX, video playback at 2x speed



RTX

PNY  
PNY PRO EXPLORE MORE.  
DELIVER MORE.

# NVIDIA RTX Server

Highly configurable reference architecture for rendering pipelines

- Powered by Quadro RTX 8000 or 6000 and NVLink
- Ray traced global illumination of up to 96 GB scenes
- Remoting, batch and multi-GPU virtualization
- Rendering time reduced from hours to minutes



Suppliers are PNY and NVIDIA authorized, like the prior VCA Certified Rendering System program

# NVIDIA RTX Server

Bare metal rendering with CUDA-X graphics

RTX Server Validation	Arnold Core Test Suite + Project SOL
Application	Autodesk Arnold 5.3.0.0
Qualified System	8x Quadro RTX 8000 or RTX 6000 + NVLink



BOXX | SUPERMICRO | EXXACT | TYAN



RTX



# NVIDIA RTX Server

Virtual desktop, batch and remoting with CUDA-X graphics

RTX Server Validation	Project Cirrus
Application	Autodesk 3DS MAX and MAYA
Remoting Protocol	Teradici Cloud Access Plus
Virtualization	Quadro vDWS v8.0 + VMWare vSphere 6.7
Qualified System	8x Quadro RTX 8000 or RTX 6000 + NVLink

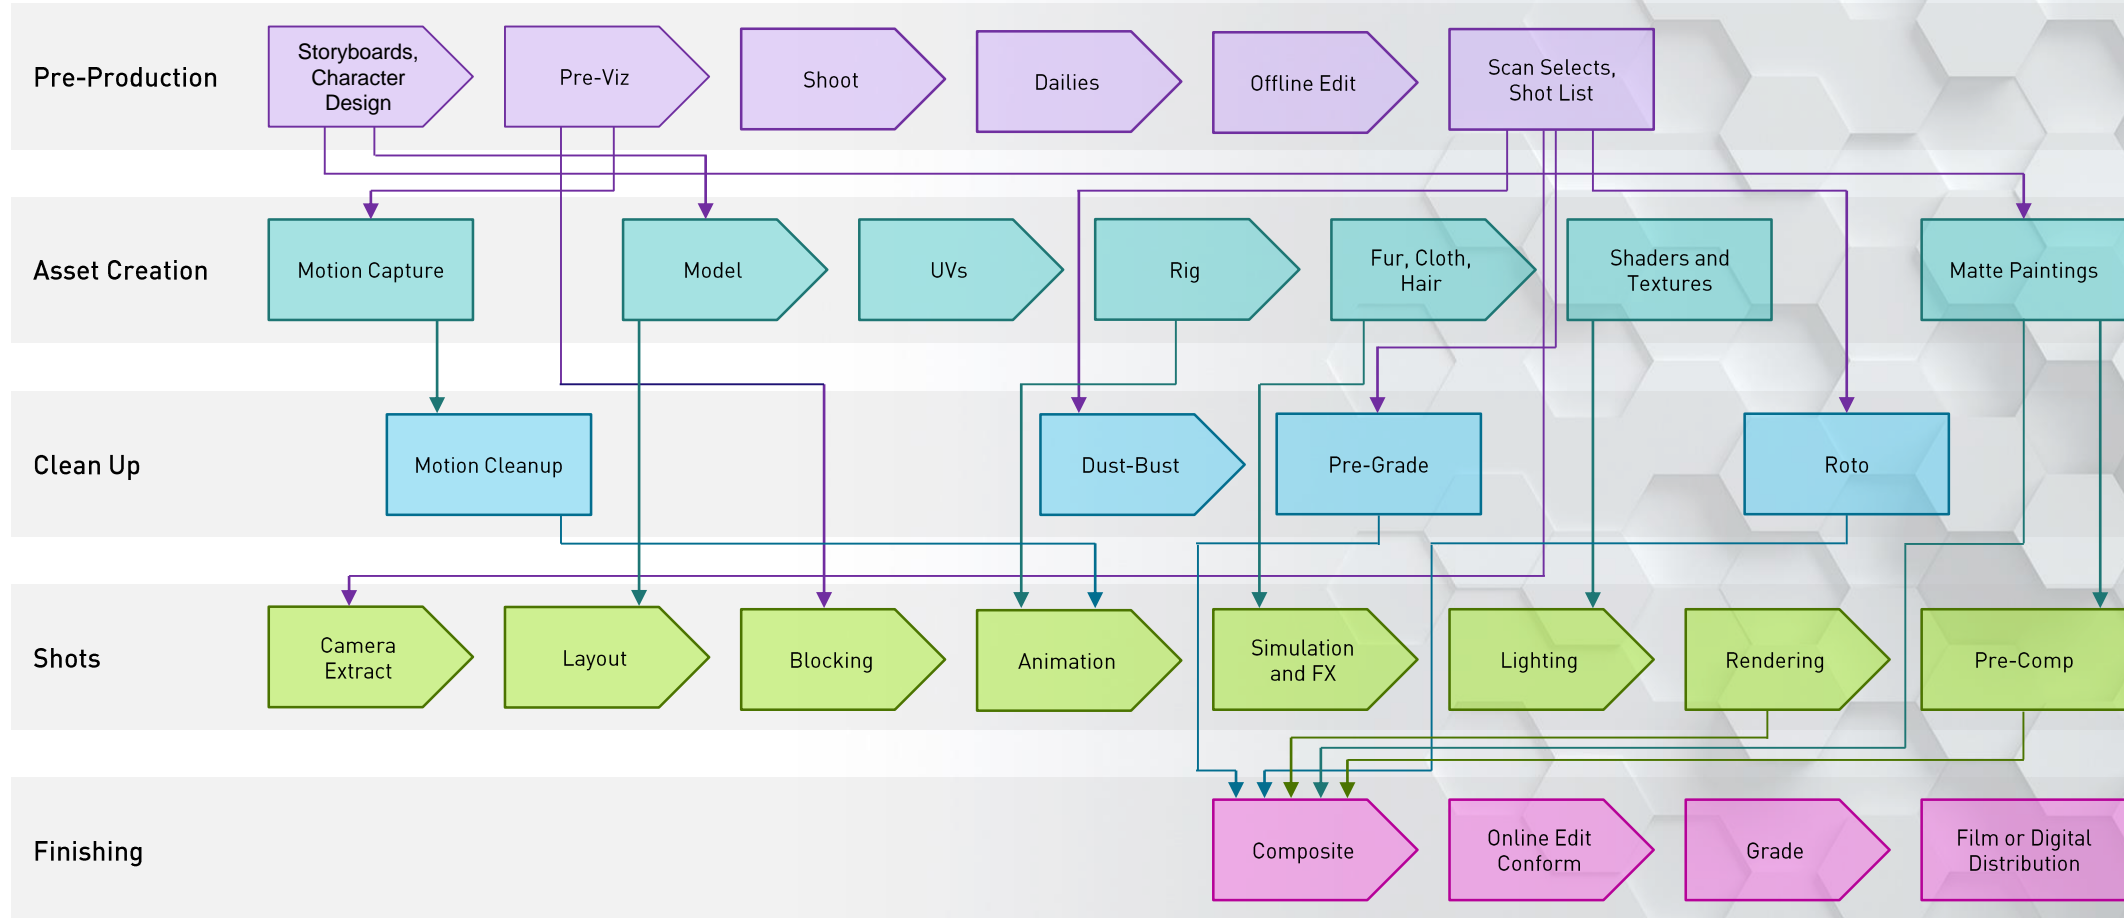


RTX



# VFX Film Pipeline

Complex, multi-phase workflow



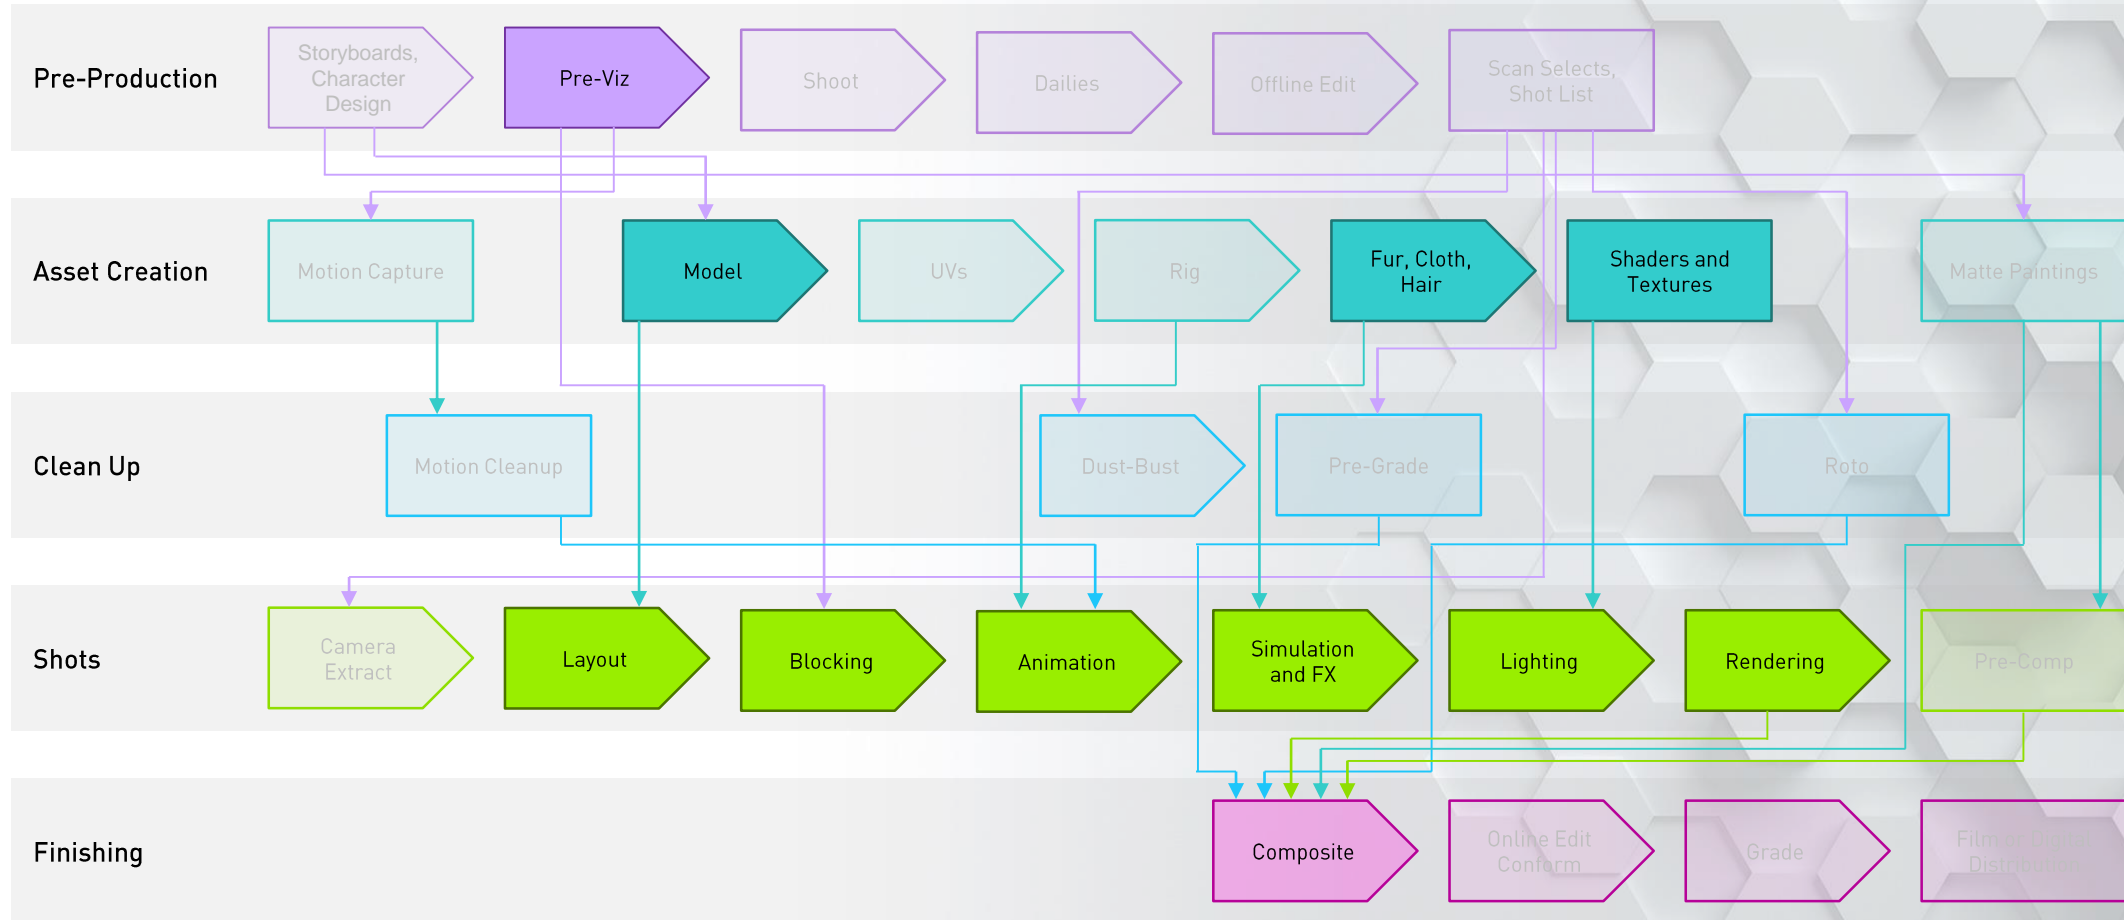
RTX



EXPLORE MORE.  
DELIVER MORE.

# VFX Film Pipeline

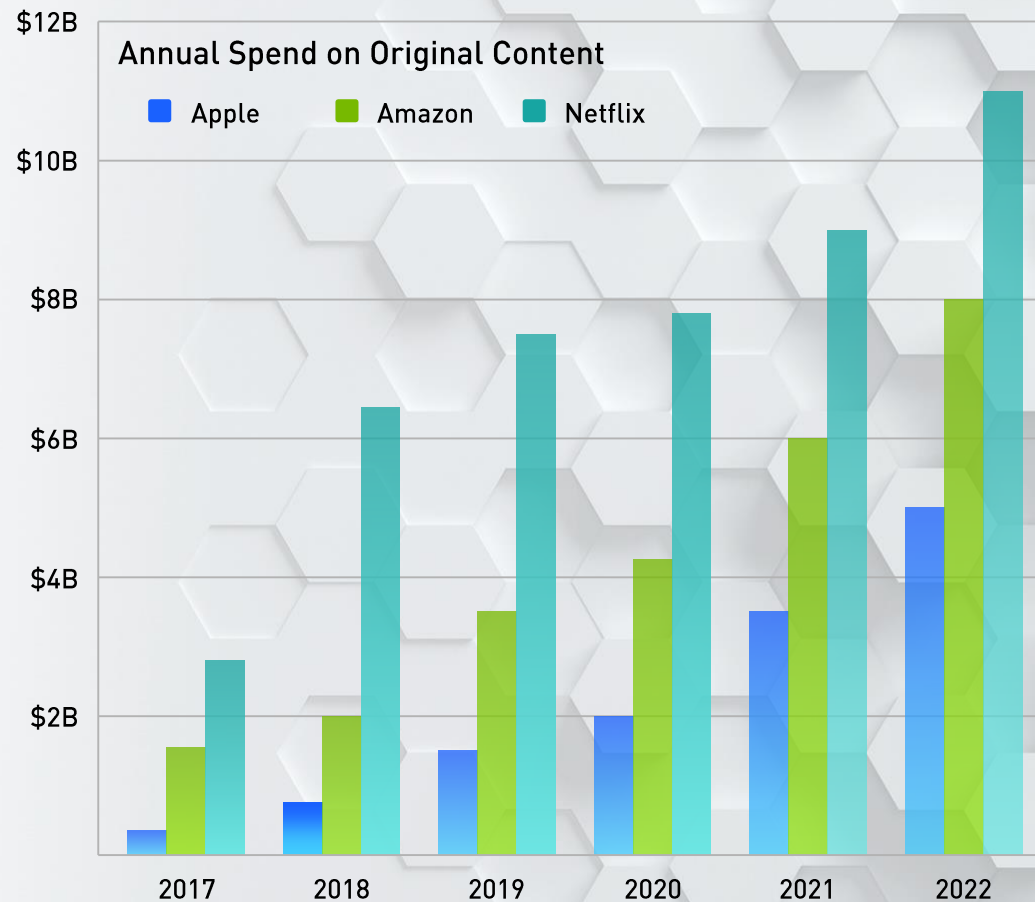
Rendering used during many phases





# NVIDIA RTX Server

## Explosion of content



# NVIDIA RTX Server

Rising quality bar requires SW and HW innovation



“After 6 months the Lightspeed and RenderMan team had a system that gives the effect of millions of lights and took the notational render time on the complex shots down from 1000 hours to 450 hours. The team continued and reduced this further to 125 hours and finally 75 hours a frame. With some additional work on the way the production team worked with the lighting in shots, the final per frame time at the end of production was just 50 hours per frame.\*

From [renderman.pixar.com](http://renderman.pixar.com)

\* Time it would take if the frame was rendered on a single core system. COCO logo, Disney, and PIXAR are registered trademarks of Disney. COCO frame grab copyright Disney - Pixar.




RTX



# NVIDIA RTX Server

## GPU memory capacity meets industry needs



24 GB

Quadro RTX 6000

24 GB

48 GB

2x Quadro RTX 6000

Quadro RTX 8000

96 GB

2x Quadro RTX 8000



RTX





EXPLORE MORE.  
DELIVER MORE.

# NVIDIA Quadro RTX Real-Time Rendering

Resets datacenter technology and economic expectations

1/4 the cost | 1/10 the space | 1/11 the power



"The more you buy, the more you save!" – Jensen Huang, CEO NVIDIA



RTX





AI


# NVIDIA NGX

A fast, easy SDK for integrating AI features

NVIDIA Develops AI Model



Create Training Data



Train and Optimize AI Model

NVIDIA Delivers AI Model and SDK for Application Integration and New Features Addition



Installs Latest AI Model with NVIDIA Driver



NGX Aware Applications Present Features

# NVIDIA NGX AI Upres

Exploring new ways to upscale content

Nearest Neighbor



Bicubic Filter



AI Super Res



AI Upres creates new pixels by interpreting the image and intelligently placing data, resulting in sharper enlargements by 2x, 4x or 8x

# NVIDIA NGX AI Slow-Mo


From 30fps to 120fps





# NVIDIA NGX AI In-Painting

A magician, not a healer





**DSW**

# Data Science Workstation

Powered by 2x NVIDIA Quadro RTX 8000 or 2x RTX 6000 with NVLink

260 TFLOPS Compute  
48 or 96 GB GPU Memory



AMAX | BOXX | COLFAX | EXXACT | Microway | THINKMATE



RTX

PNY  
PNY PRO EXPLORE MORE.  
DELIVER MORE.

# NVIDIA CUDA-X AI and NVIDIA RAPIDS

Executes end-to-end data science and analytics pipelines entirely on RTX

## Hassle-Free Integration

Accelerate your Python data science toolchain with minimal code changes and no new tools to learn

## Improves Model Accuracy

Increase machine learning model accuracy by iterating on models faster and deploying them more frequently

## Reduces Training Time

Drastically improve your productivity with near-interactive data science

## Open Source, NVIDIA Optimizations

Customizable, extensible, interoperable open-source software is optimized and supported by NVIDIA and built on Apache Arrow

RAPIDS



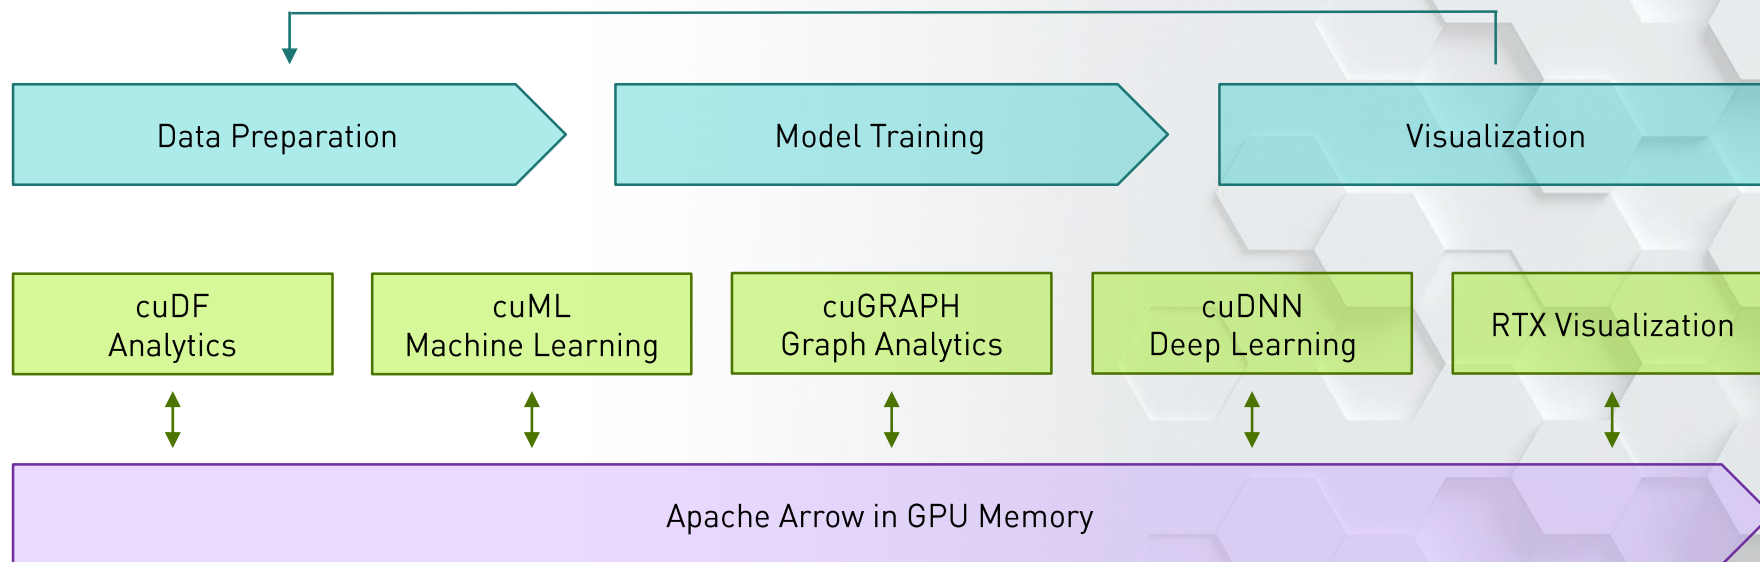
RTX



EXPLORE MORE.  
DELIVER MORE.

# NVIDIA RAPIDS with Anaconda

## The new GPU data science pipeline



### Apache Arrow

A columnar in-memory data structure that delivers efficient and fast data interchange with flexibility to support complex data models

### cuDF Analytics

cuDF is a DataFrame manipulation library that accelerates loading, filtering, and manipulation of data for model training data preparation

### cuML Machine Learning

cuML provides GPU-accelerated versions of all machine learning algorithms available in scikit-learn

### cuGRAPH

A framework and collection of graph analytics libraries that seamlessly integrate into the RAPIDS data science platform

### cuDNN

RAPIDS provides native array\_interface support, so data can be pushed to DL frameworks like PyTorch and Chainer

### RTX Visualization

Native GPU in-memory data format provides high-performance, high-FPS data visualization, even with very large datasets



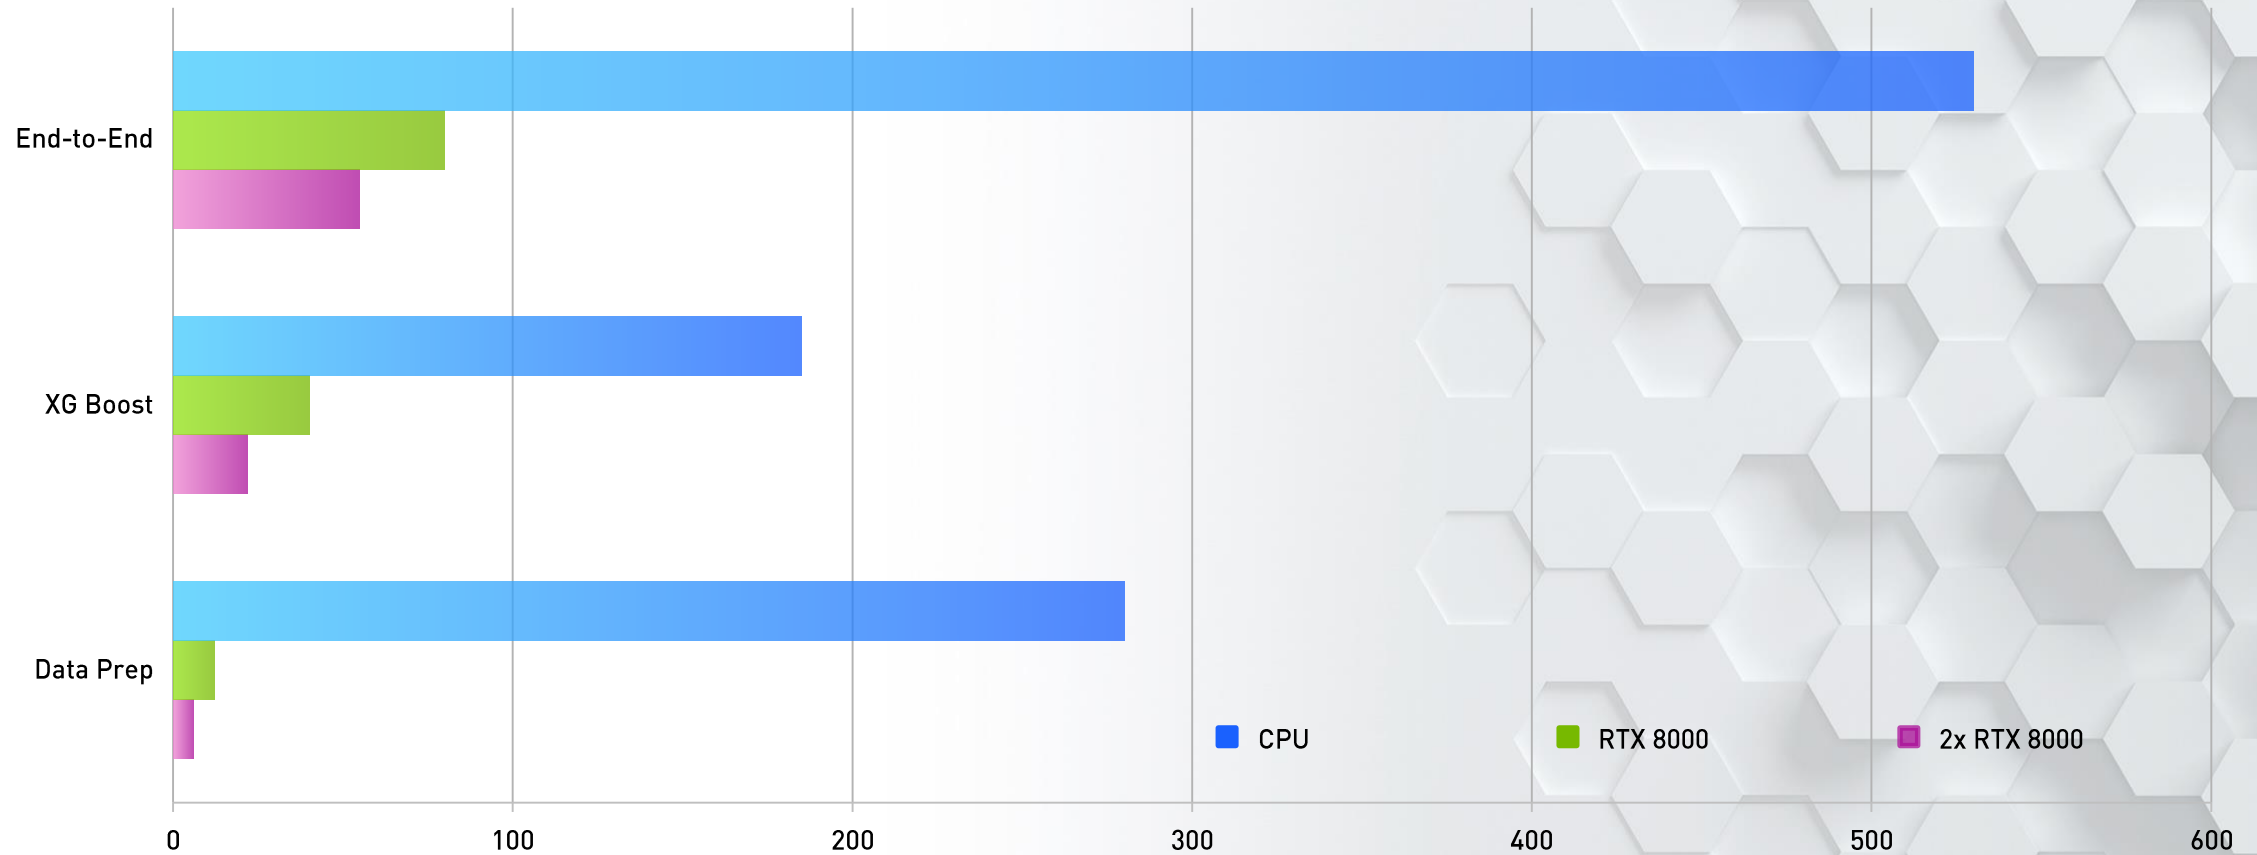
RTX



EXPLORE MORE.  
DELIVER MORE.

# RAPIDS with Quadro RTX 8000

Unprecedented data science performance



\*End-to-End time = ETL + conversion + training + validation. CPU Xeon 6140 at 3.2 GHz, 3.7 GHz Turbo, 384 GB RAM, Ubuntu 16.04.4, NVIDIA driver 410.93



RTX

PNY  
PNY PRO

EXPLORE MORE.  
DELIVER MORE.

# Data Science Workstation

Opening up new vistas of discovery



**RTX**

**PNY**  
EXPLORE **MORE.**  
DELIVER **MORE.**  
**PNY PRO**



# NVIDIA Quadro RTX

The Fusion of Graphics and AI

S9969 | Booth 1133



**RTX**



EXPLORE **MORE.**  
DELIVER **MORE.**