Secure Content Creation in the Cloud

Rick Champagne

Global Media & Entertainment Industry Management, NVIDIA

Ian Main

Technical Principal, Teradici

AUTODESK. UNIVERSITY

EXPLOSION OF CONTENT

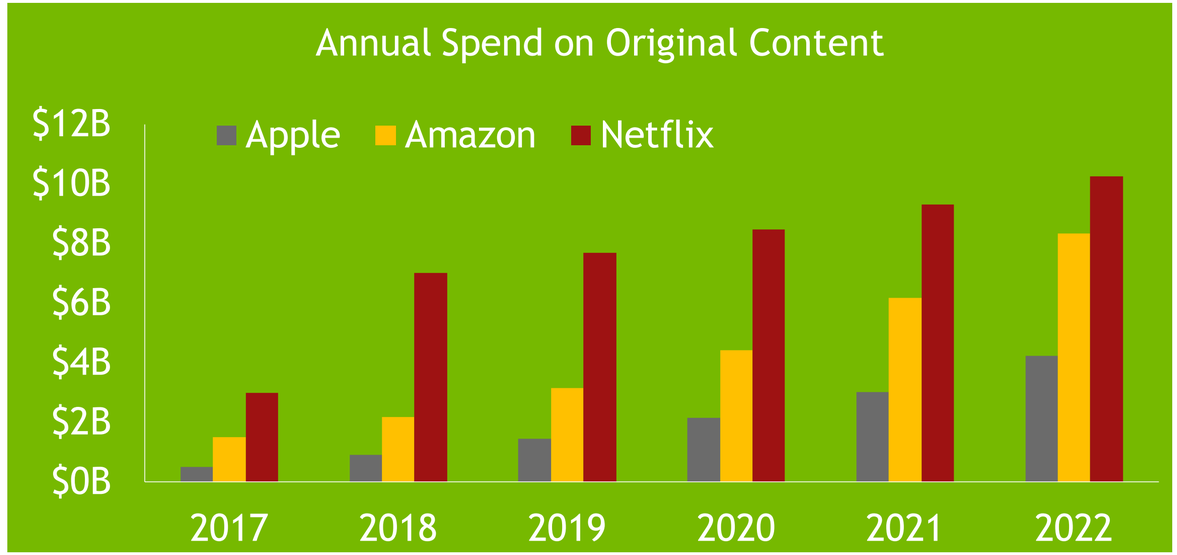

RISING QUALITY & COMPLEXITY

GLOBALIZATION

CONTENT SECURITY

NVIDIA.

NVIDIA Turing Architecture

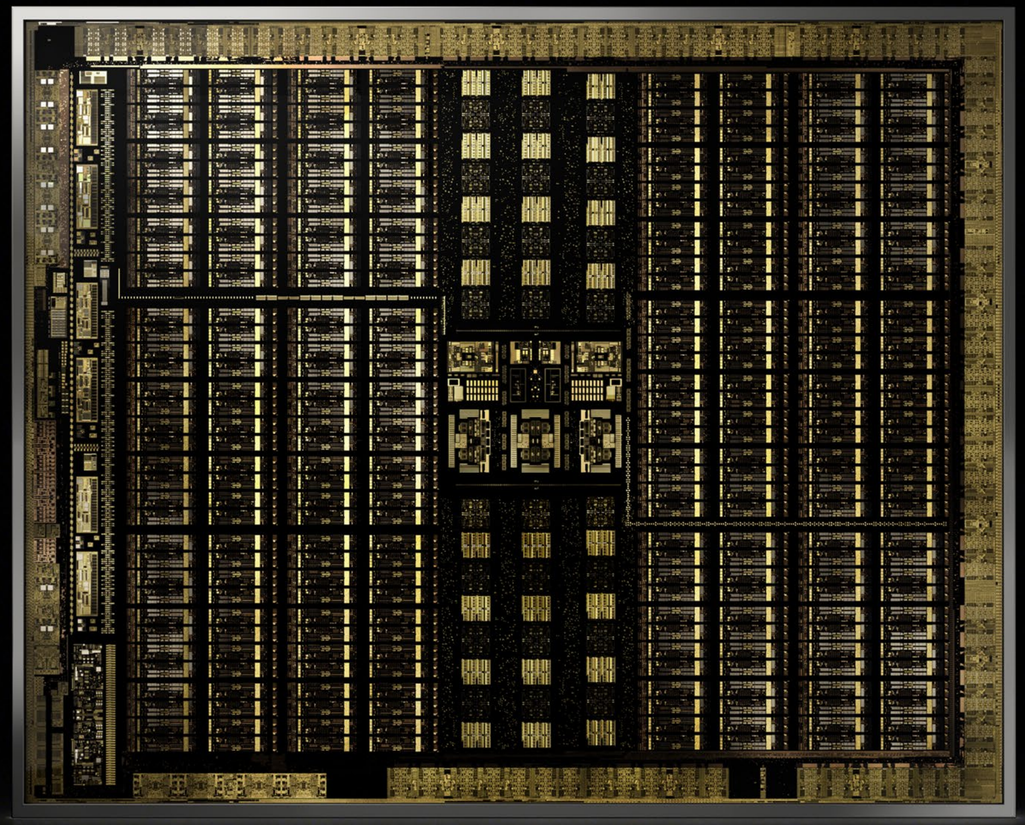

Turing SM
16 TFLOPS + 16 TIPS
Concurrent FP & INT Execution
Unified L1 Cache
Variable Rate Shading

RT Core
10 Giga Rays/sec
Ray Triangle Intersection
BVH Traversal

Tensor Core
125 TFLOPS FP16
250 TOPS INT8
500 TOPS INT4

Display
Native HDR
8K DisplayPort
VirtualLink

NVLINK
100 GB/sec
GPU-GPU Memory Access

Video
HEVC 8K Real Time Encode
25% Improved Bitrate

Memory
6MB L2 Cache
384-bit G6 @ 14Gbps
672 GB/sec

NVIDIA

NVIDIA RTX PLATFORM

New Generation Of Hybrid Rendering

- Interoperability Between Rasterization, Ray Tracing, Compute, AI
- New Turing Ray Tracing Acceleration in OptiX, DXR, Vulkan
- New NVIDIA MDL Materials Open Source
- New Support for Pixar Universal Scene Description (USD)

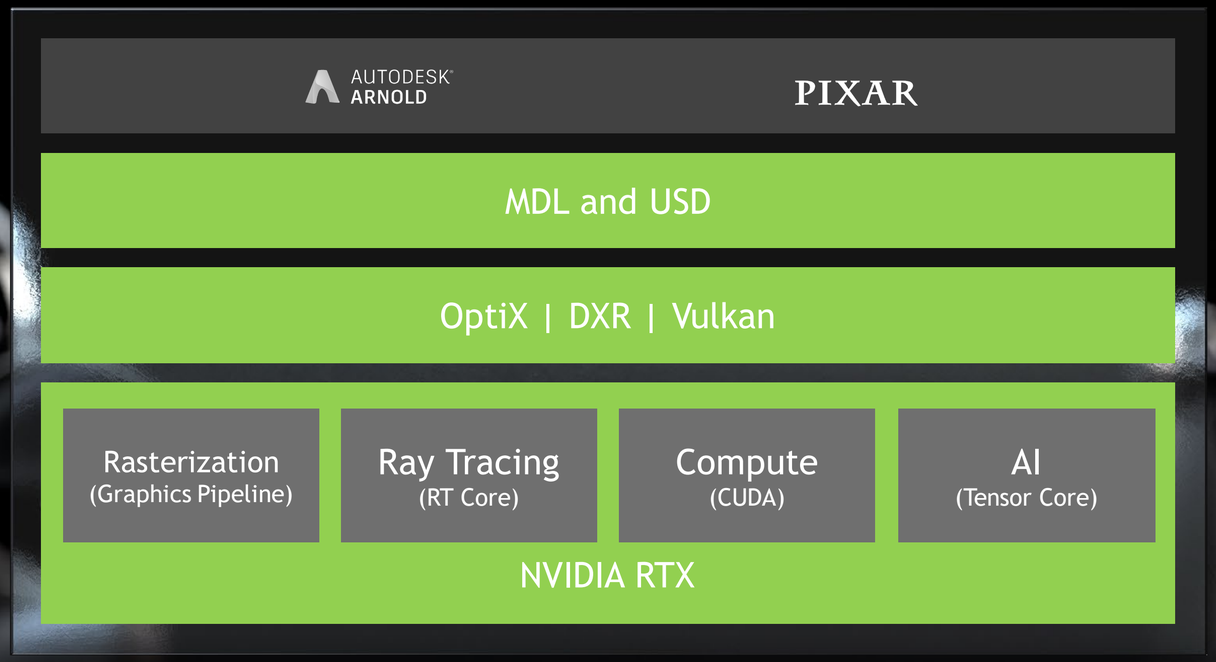

NVIDIA

Turing Ray Tracing Performance

>10 Giga Rays

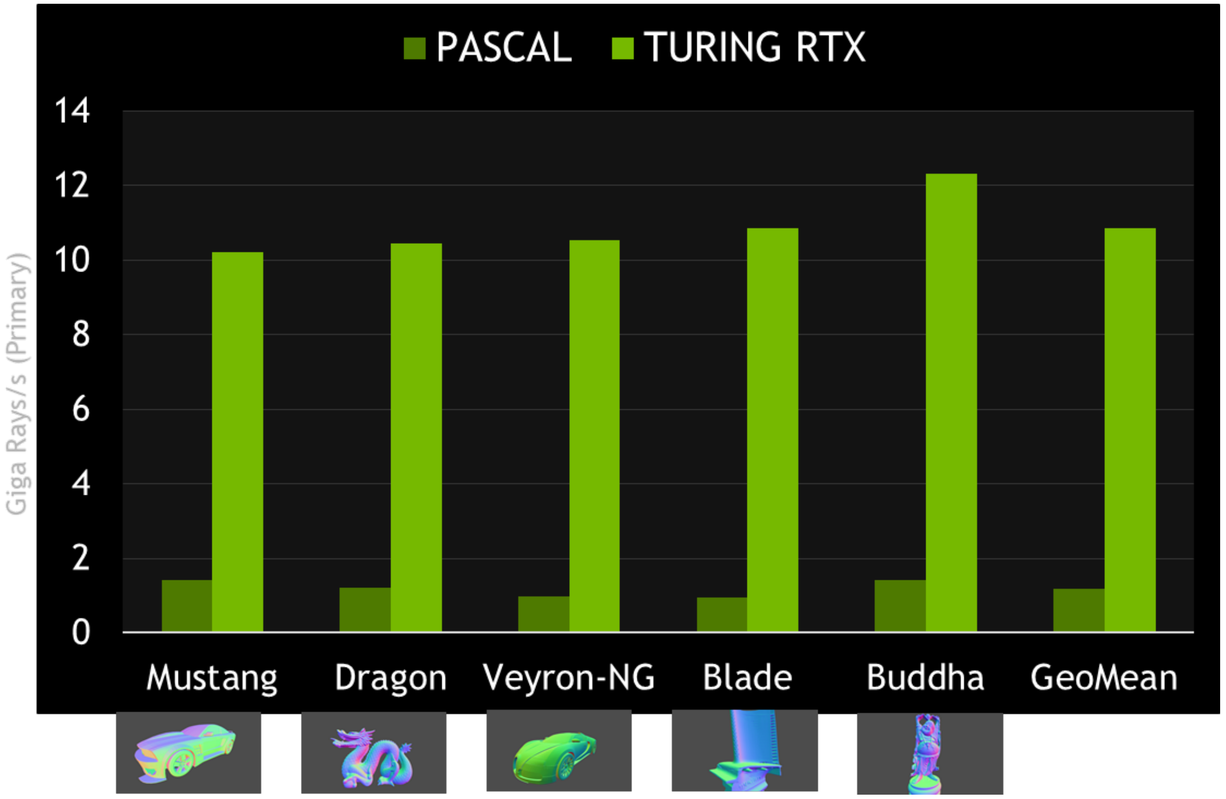

PASCAL	TURING
11.3 TFLOPS	68 RT Cores
1.1 Giga Rays	10+ Giga Rays
10 TFLOPS / Giga Ray	~10X faster than PASCAL

NVIDIA RTX Server

Flexible, highly configurable solution:

Configured with Quadro RTX 8000 or RTX 6000 GPUs

Up to 96 GB of ultra-fast memory per GPU-pair with NVLink

Certified for Autodesk Arnold GPU

NVIDIA

NVIDIA RTX Server

Exponential Power at a Fraction of the Cost

240 Dual 12-core Skylake CPU Servers
144 kW
$2M Render Farm

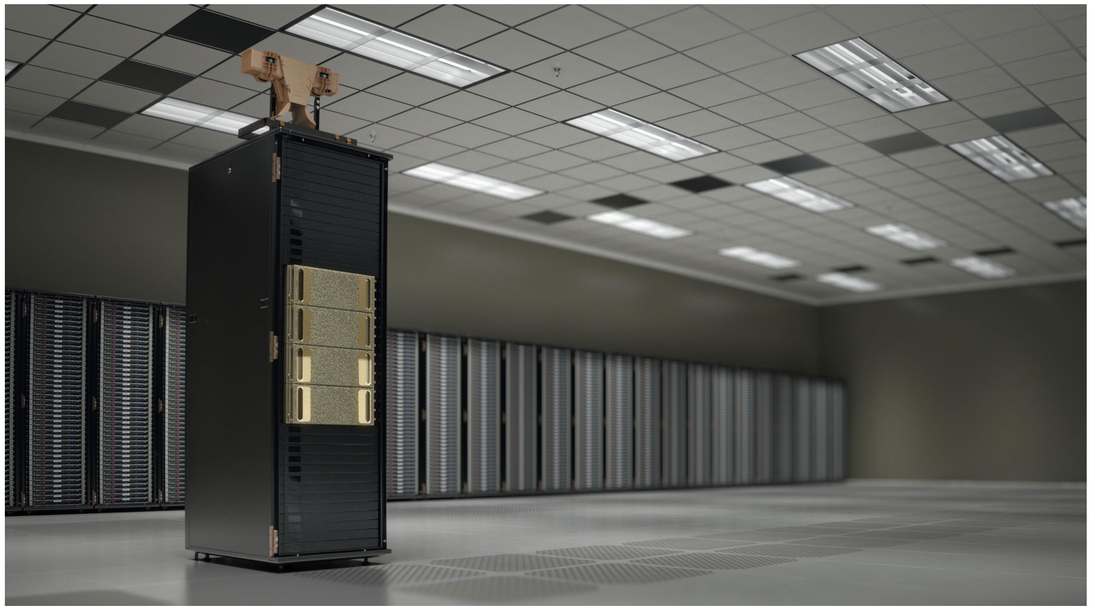

4 RTX 8-GPU Servers
13 kW
$500,000

1/4 the Cost **1/10 the Space** **1/11 the Power**

Interactive Content Creation | **Batch Rendering**

RTX-Accelerated Rendering

RTX SERVER VALUE ACROSS USE CASES

Offline/Batch Rendering

Accelerate offline batch rendering with the power of GPU acceleration by connecting one or more workstations to NVIDIA RTX Server.

Accelerated Desktop Rendering

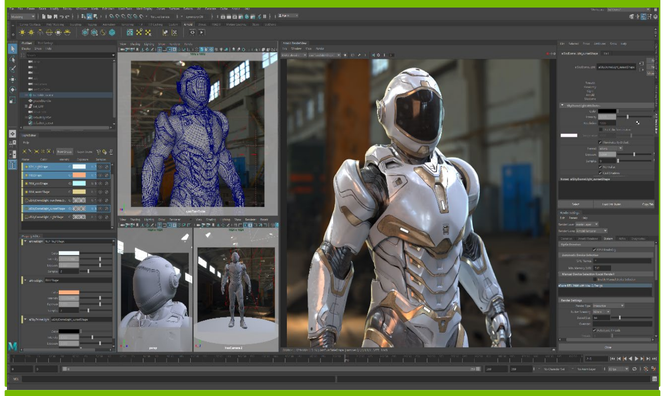

Boost local workstation rendering performance and interactivity by connecting to NVIDIA RTX Server and tapping into the power of up to 10 GPUs.

Virtual Workstations

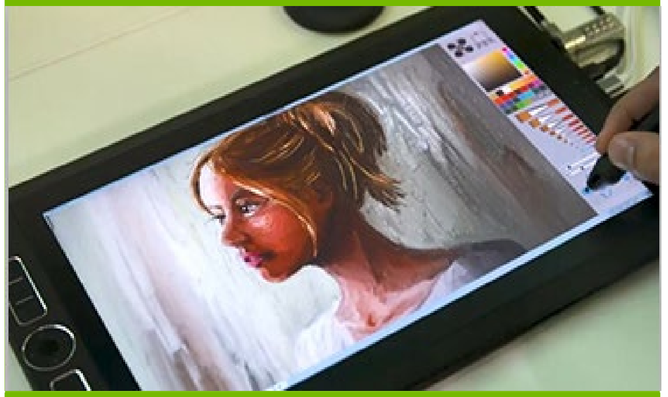

Provision multiple virtual workstations or a combination of virtual workstations and render nodes from a single NVIDIA RTX Server with Quadro vDWS.

“We never expected to see results this dramatic. This will completely change how our artists work.”

- Michele Sciolette, CTO, Cinesite

AUTODESK. CINESITE STUDIOS

2 Shots a Day

7 Shots a Day

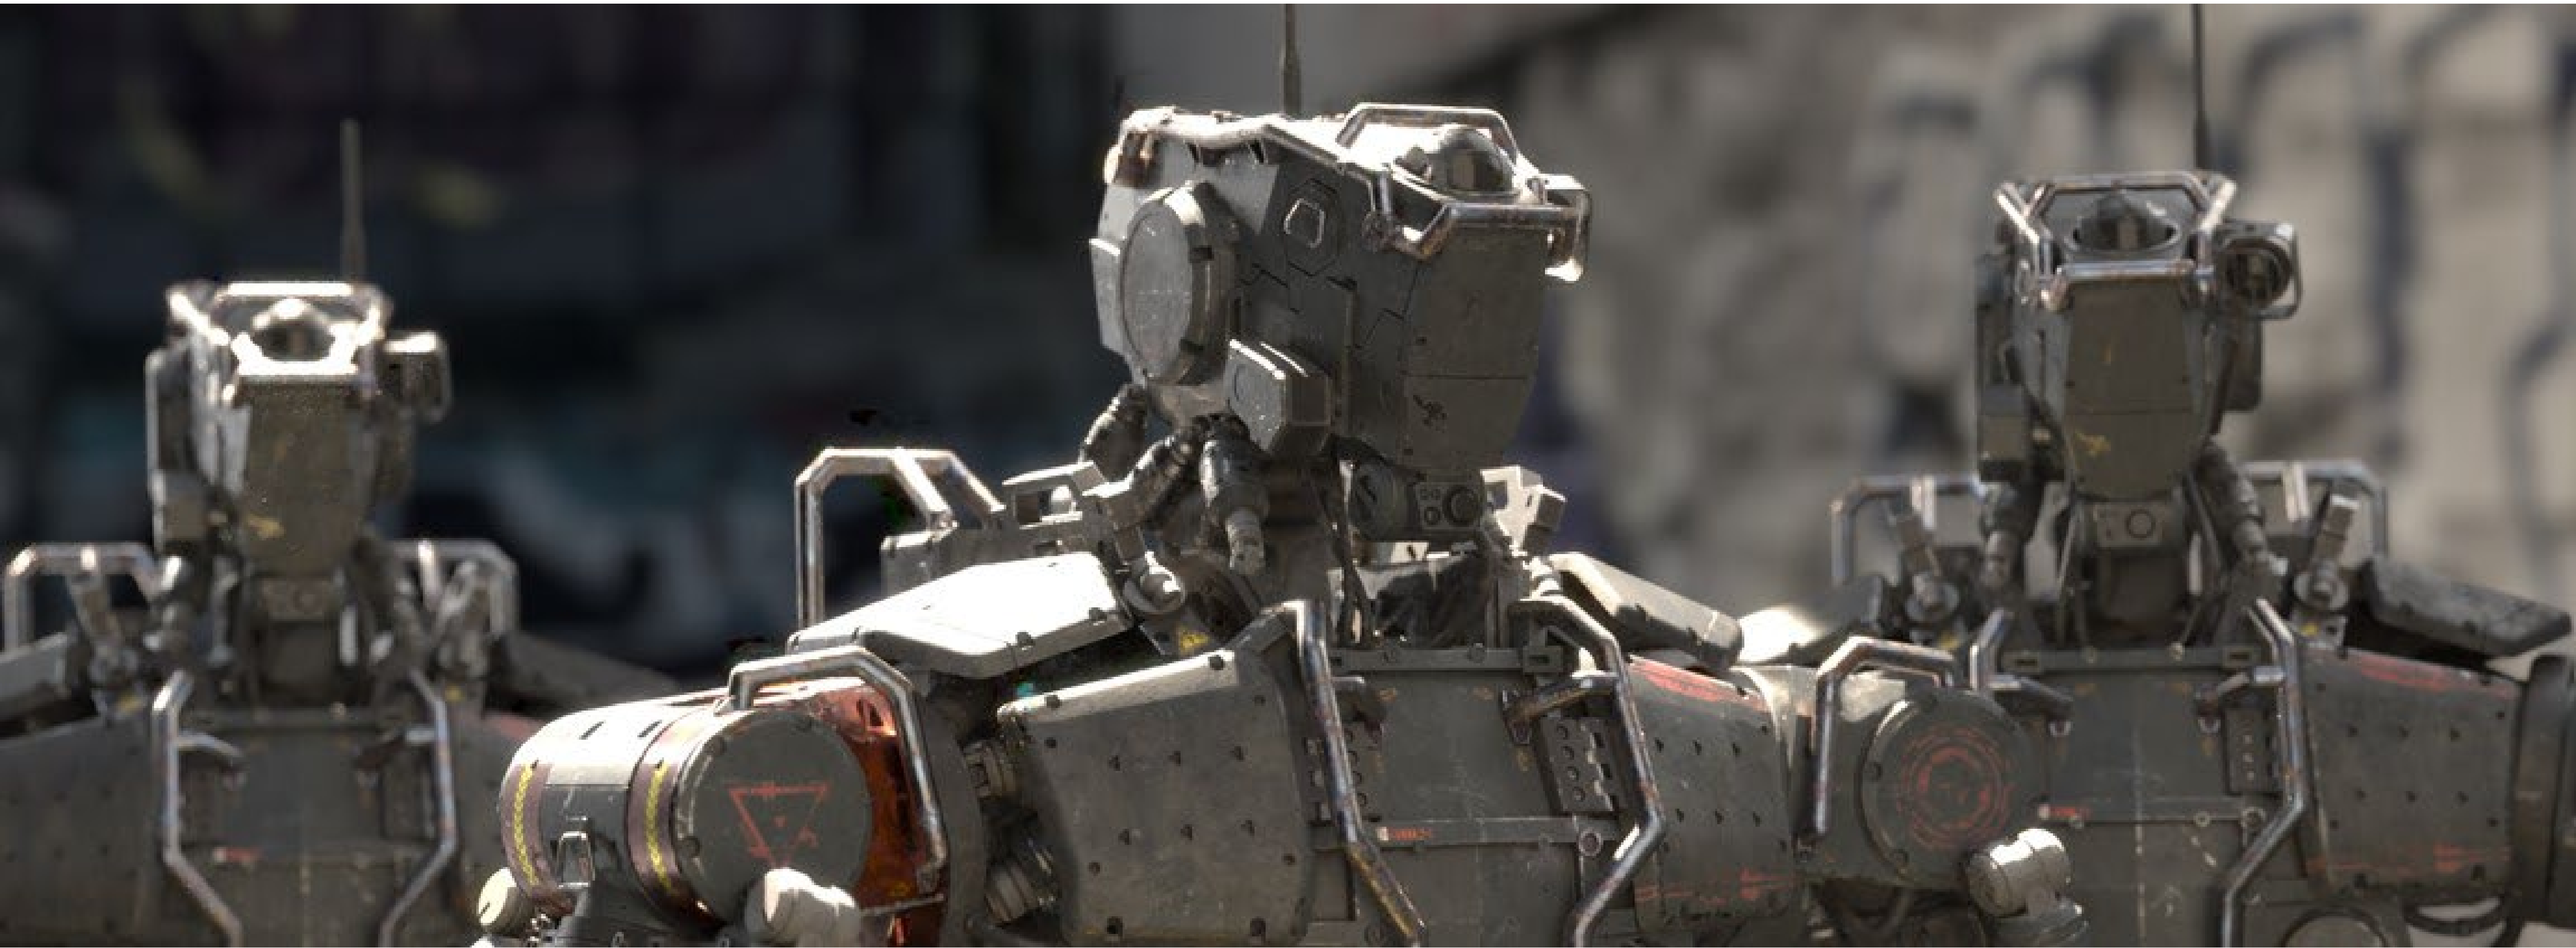

IMAGE ENGINE

Image Engine shows RTX Server boosting production workflows using graphics from Netflix's Lost in Space.

" We are thrilled to collaborate with NVIDIA and Autodesk on the next generation of high-performance production rendering. Delivering the best creative results fast is critical to our success. GPU rendering on NVIDIA RTX Server with Autodesk Arnold promises to deliver the exponential leap in speed we hope to utilize in all production departments."

CARSTEN KOLVE
Digital Supervisor | Image Engine

image engine
VISUAL EFFECTS + ANIMATION

	# of Nodes	Total Render Time	Cost of Power (5 yrs)	Total Cost
CPU - Dual Skylake	25	38 Hours	$70k	$250k
RTX Server - 4x RTX 8000	1	6 Hours	$10k	$30k

RTX SERVER WITH VIRTUAL WORKSTATIONS

INCREASE PRODUCTIVITY

Get up to 5X faster rendering per 1U versus a CPU-only solution and scale up to 4U to maximize rendering speed by 20X.

IMPROVE UTILIZATION

Consolidate multiple physical workstations to one server with multiple virtual workstations, and suspend workstation nodes/spin up as needed

REDUCE COSTS

Buy one virtualized server instead of multiple CPU-only servers, which saves money and is easier to manage.

NVIDIA T4 GPU

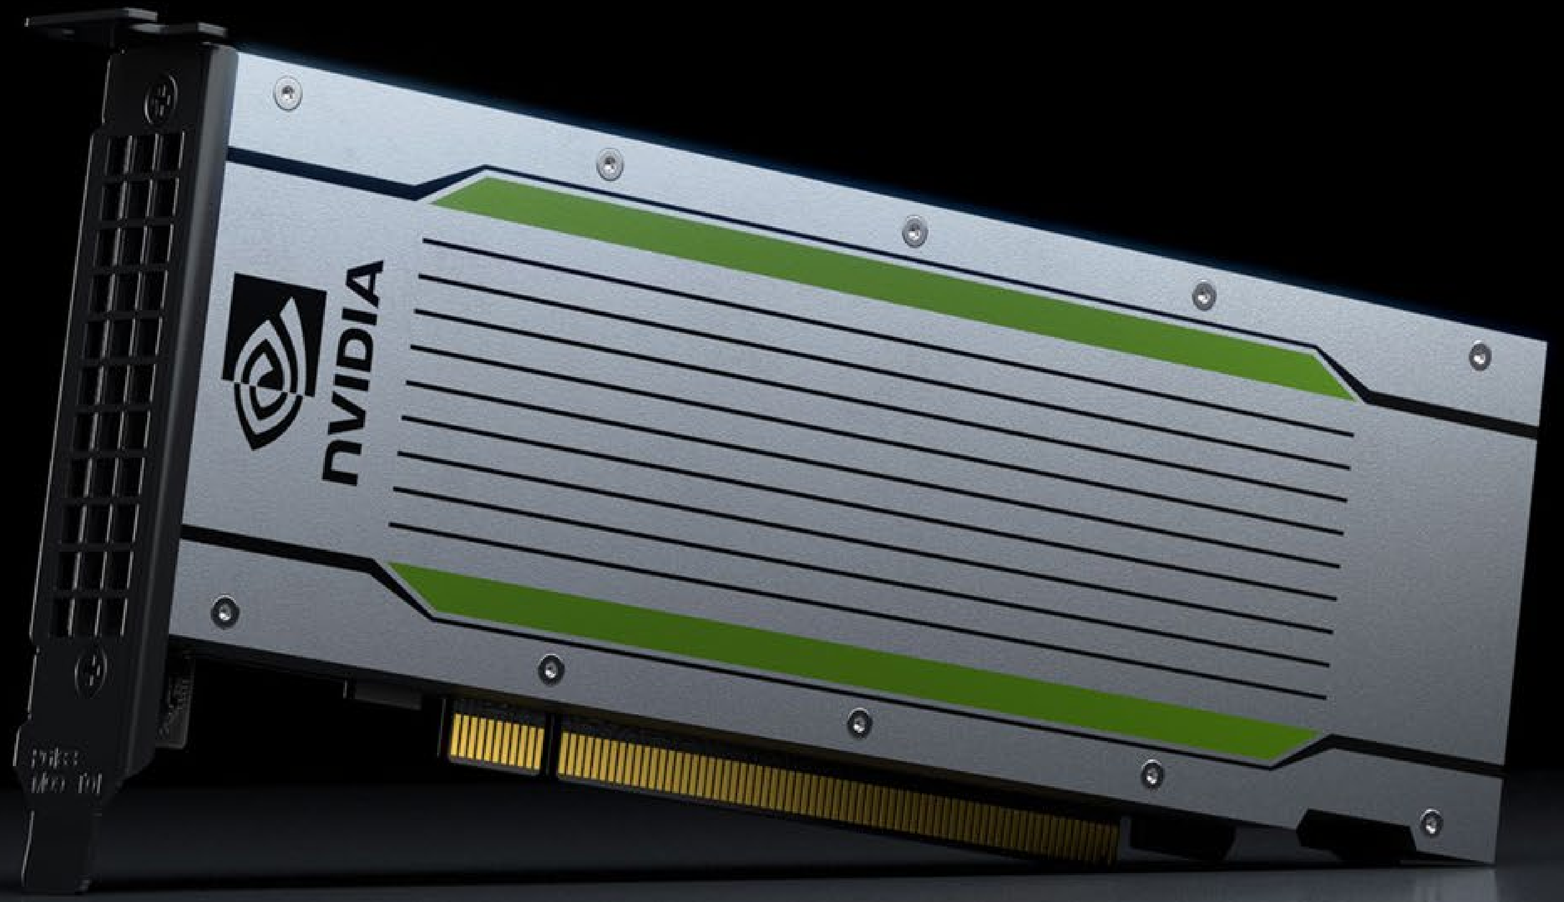

CUDA Cores	2560
Turing Tensor Cores	320
RT Cores	40
Memory BW	320+ GB/s
Frequency	1.005 GHz
TFLOPS (FP32)	7.7
Memory	16 GB

NVIDIA Virtual GPU for Enterprise

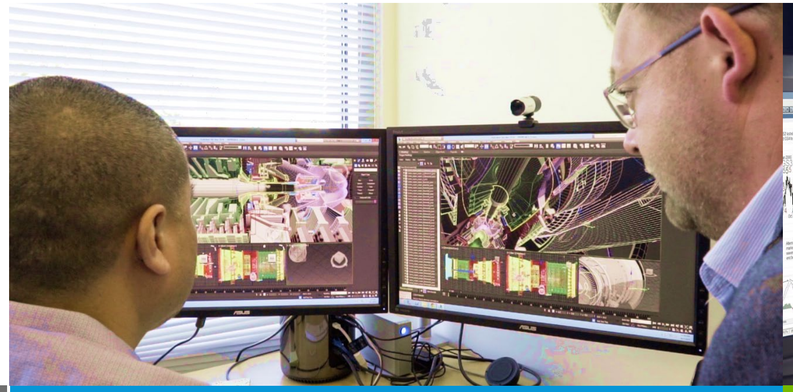

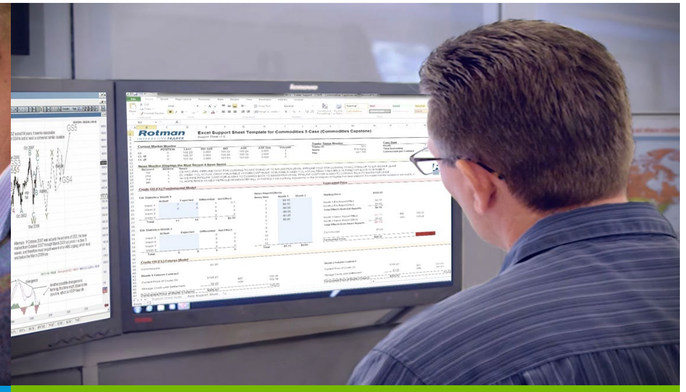

Type of User	**Creative/Technical Professional**	**Knowledge Worker**
Virtual GPU Software Edition	**Quadro Virtual Data Center Workstation**	**GRID Virtual PC/Virtual Applications**
GPU Hardware Recommendation	NVIDIA T4 Tesla P40, RTX 6000, RTX 8000, V100 Tesla P6 for blade form factor	Tesla M10, NVIDIA T4 Tesla P6 for blade form factor

Overcome the Challenges of Mobility and Collaboration in M&E

Increase Productivity

Deliver superior graphics performance on virtual desktops for a satisfying user experience.

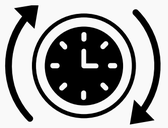

Collaborate Anywhere, on Any Device

Speed the creative process by giving employees the freedom to access fully 3D-capable virtual workstations on any device.

Uncompromised Security

Onboard new contractors in minutes while ensuring the security of media assets in the data center or cloud.

Driving Data Consistency

Centralize projects in the data center or cloud for greater consistency and control over changes and quality.

Who is Teradici?

PCoIP Protocol

Lossless, color-accurate, highly-responsive, enhanced with PCoIP Ultra™ features

13+ Million

PCoIP users worldwide and growing

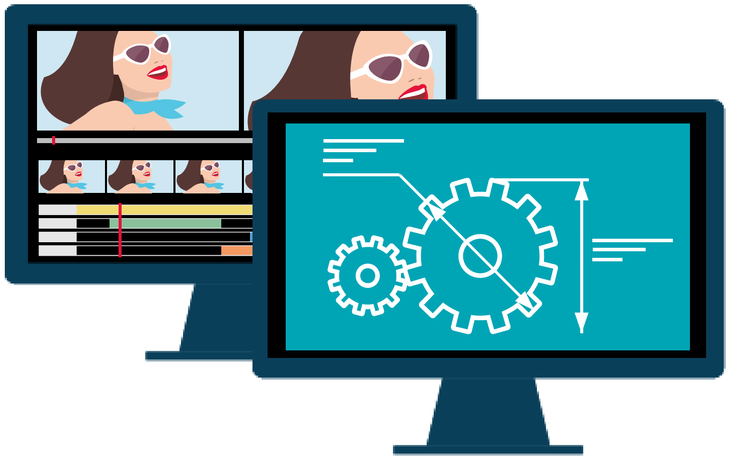

Industries

M&E, Manufacturing, Design & Engineering, Healthcare, Finance, Government, etc.

What is the PCoIP® Protocol?

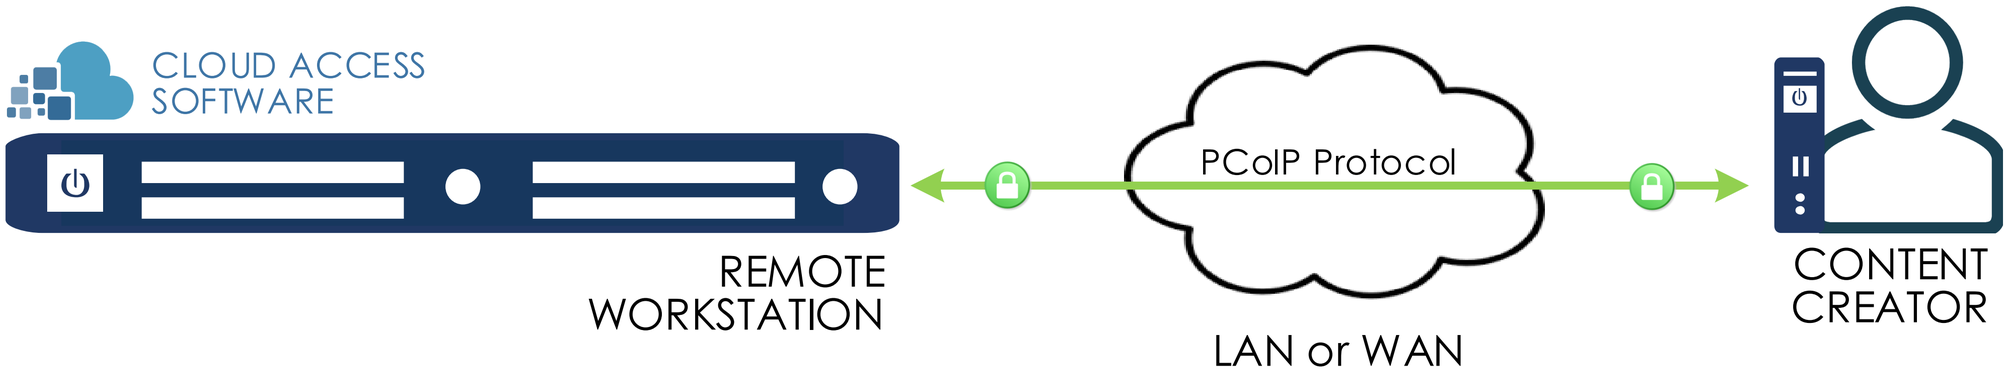

Baseline LAN Deployment

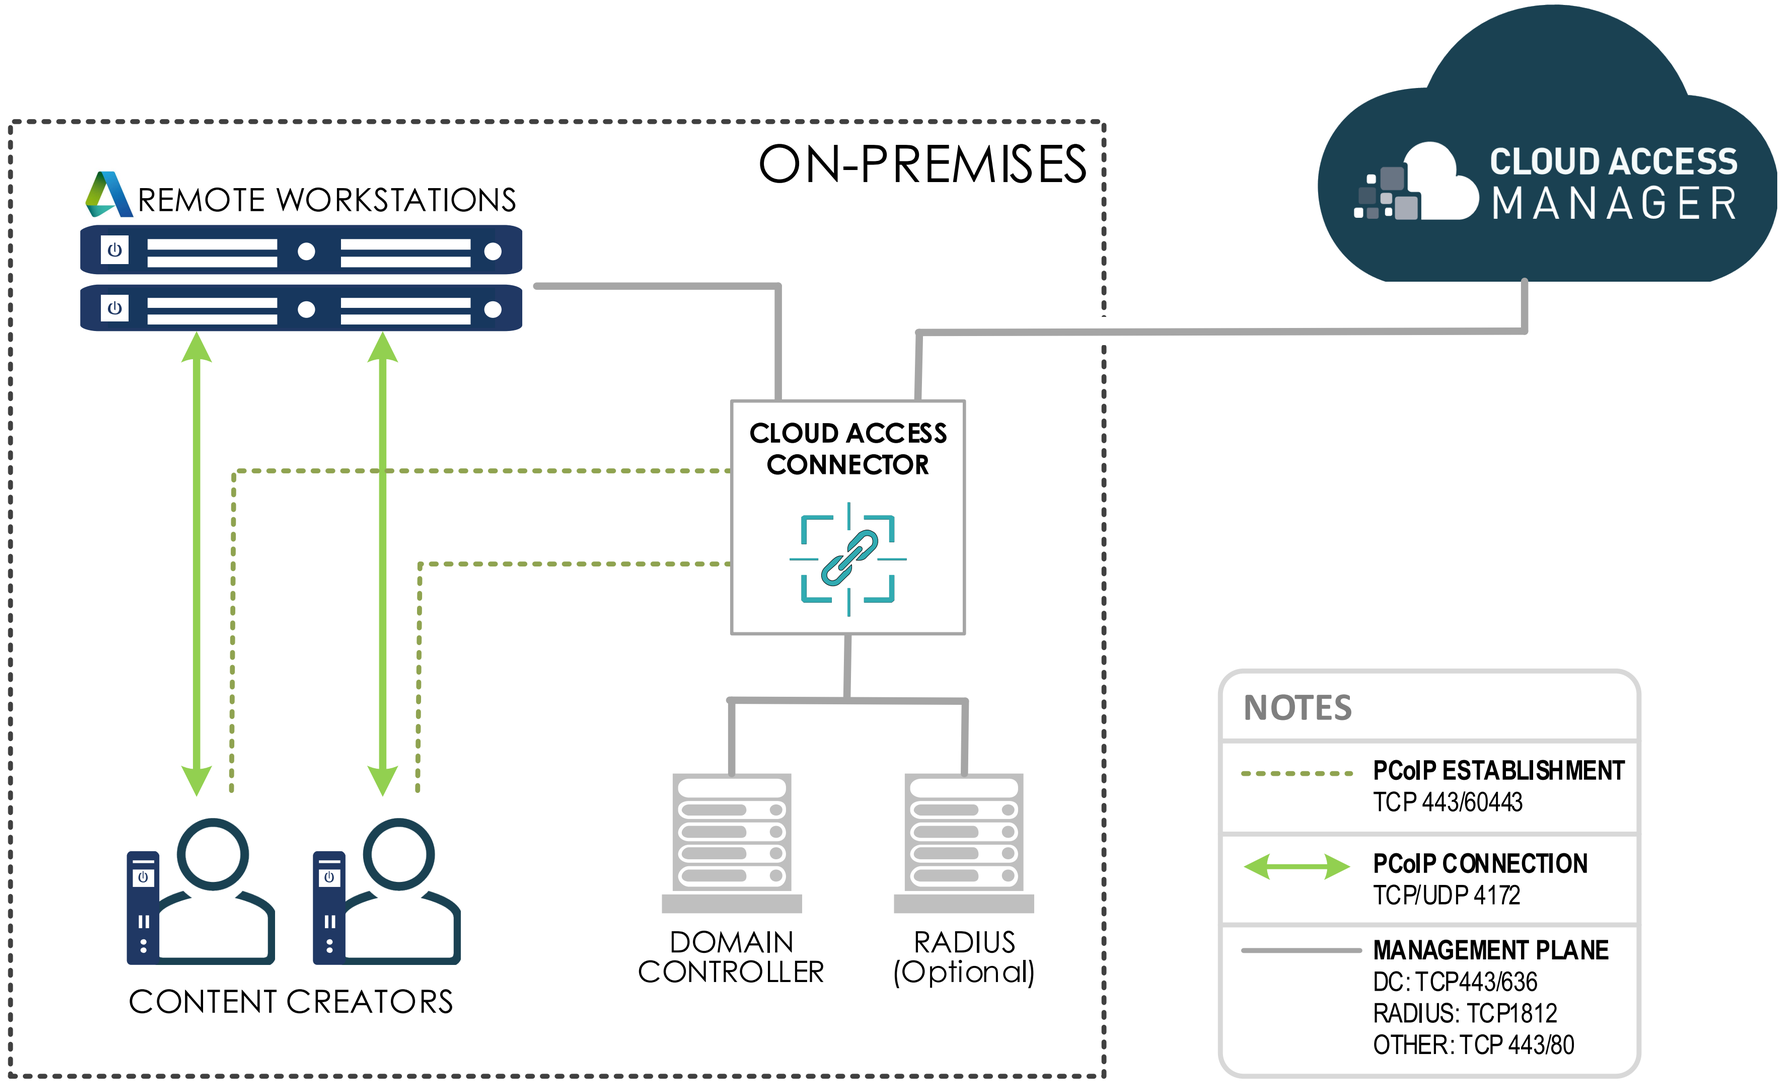

External Access to On-Premises Resources

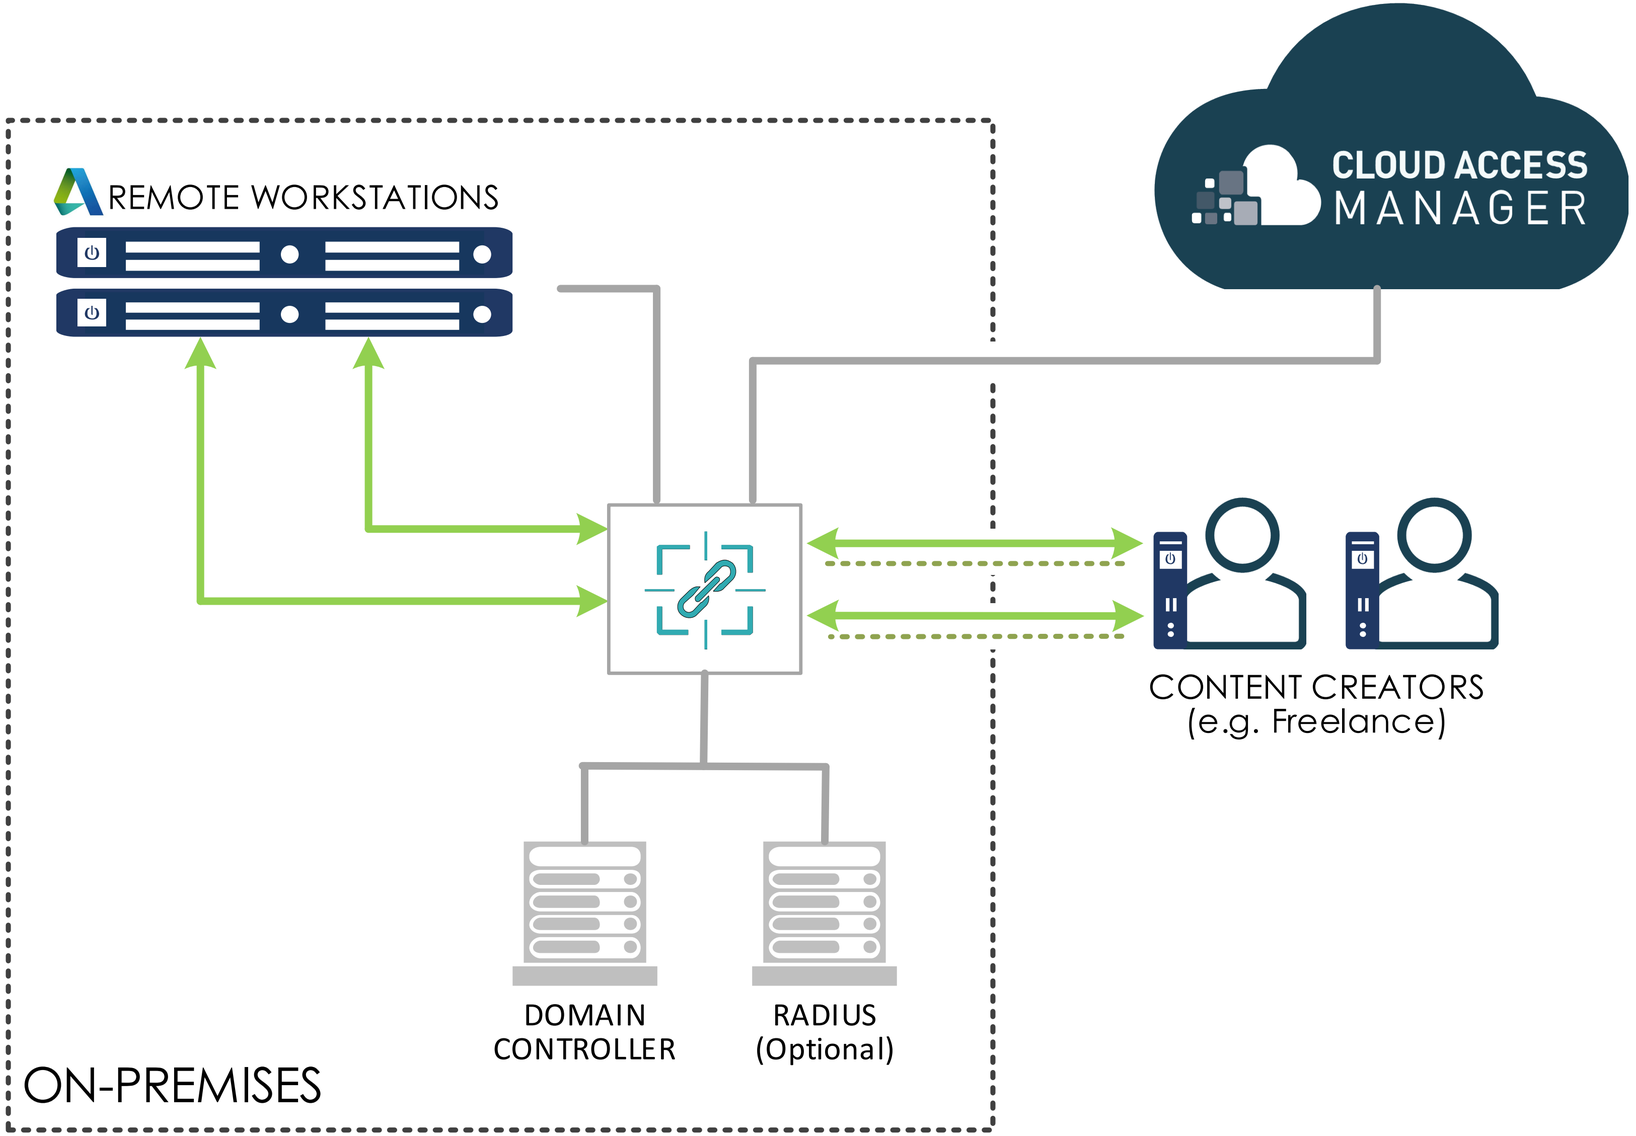

Public Cloud Deployment with On-Site or External Users

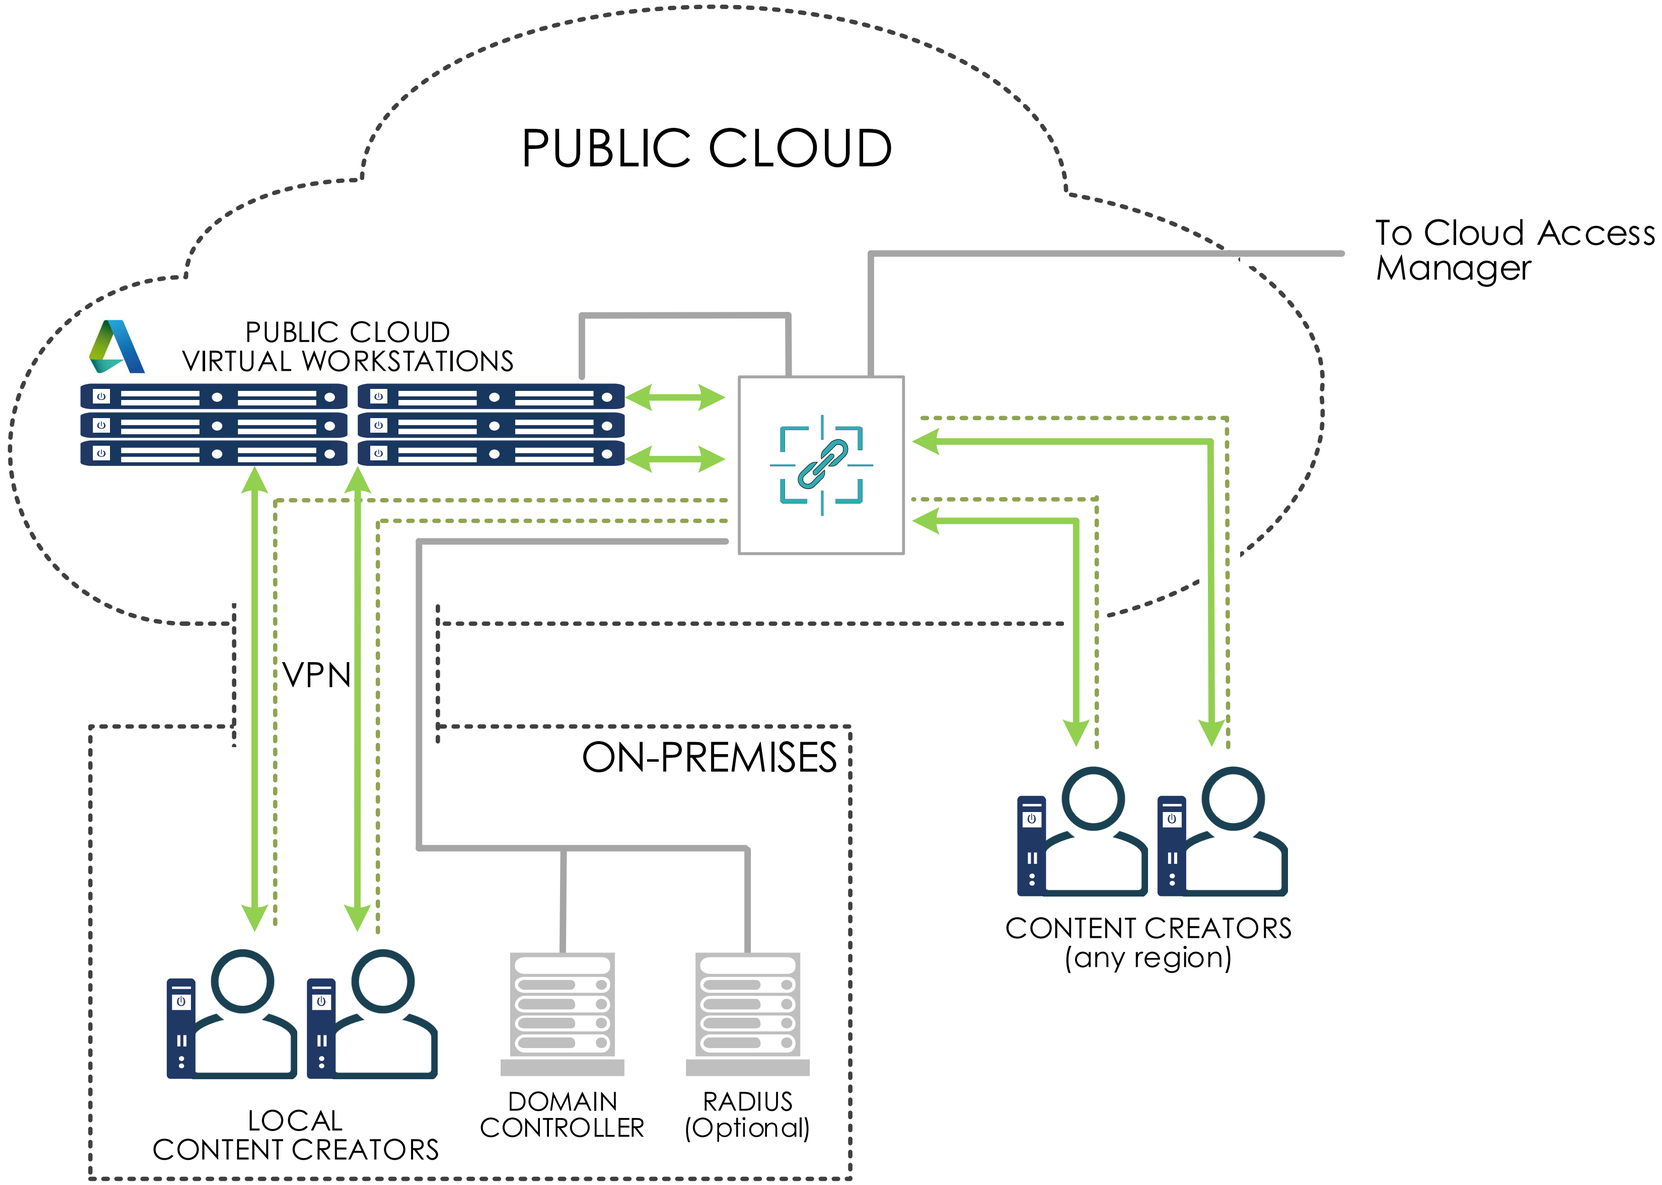

Multicloud Deployment

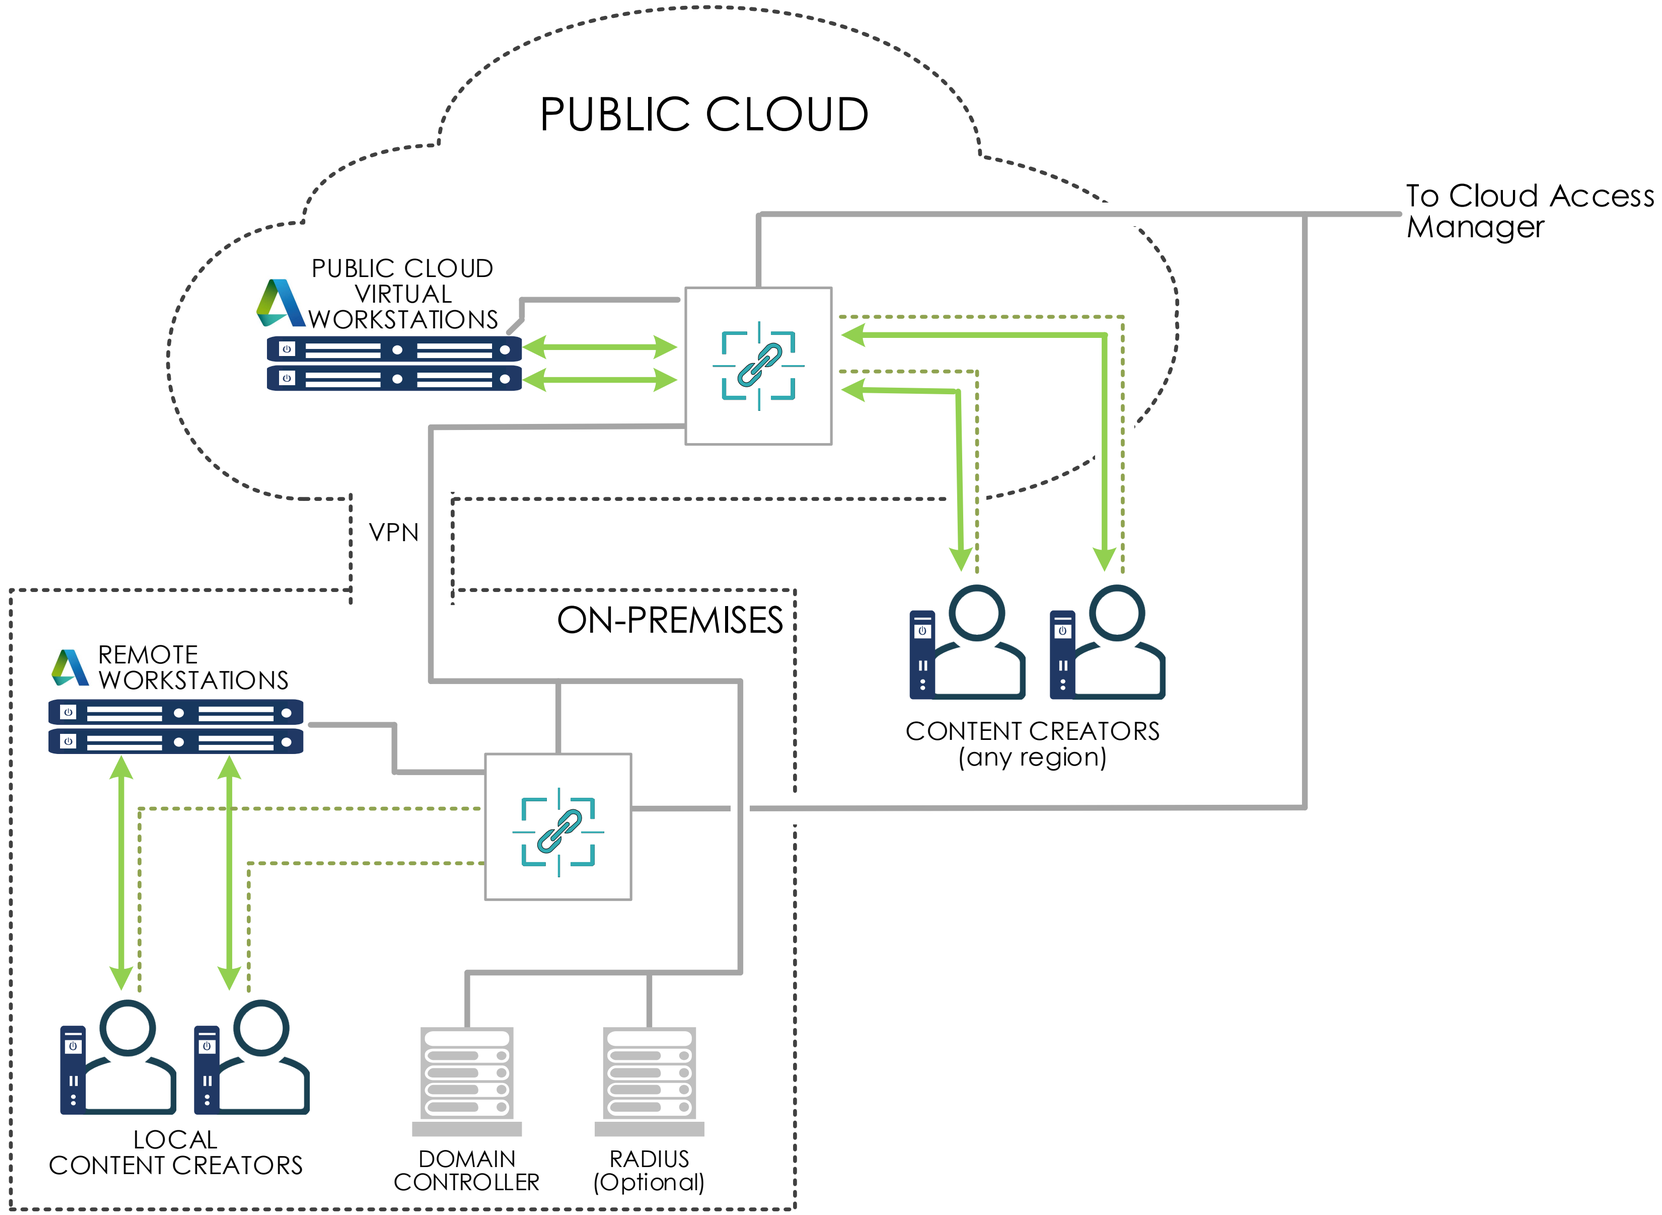

Public Cloud Support for Virtual Workstations

Cloud Service Provider	NVIDIA Quadro VDWS Virtual Workstation Offering
aws	T4, V100, K80, M60
Azure	“NV Series”: M60, “NC Series”: P100
Google Cloud	T4, V100, K80, P4, P100

Choosing a GPU

	NVIDIA Tesla K80	NVIDIA Tesla P100	NVIDIA Tesla V100	NVIDIA Tesla P4	NVIDIA Tesla T4
CUDA Cores	2496	3584	5120	2560	2560
Memory BW	240 GB/s	732 GB/s	900 GB/s	192 GB/s	320 GB/s
Frequency	0.562 GHz	1.13 GHz	1.53 GHz	0.810 GHz	1.005 GHz
TFLOPS (FP32)	2.91	5.3	14	5.5	7.7
Memory (GB)	12	16	16	8	16
Config	1, 2, 4, 8	1, 2, 4	1, 2, 4, 8	1, 2, 4	1, 2, 4
Visualization	No	Yes	No	Yes	Yes
Price/hr (USD)*	$0.45	$1.46	$2.48	$0.60	$0.96

** Price per hour, in USD, in us-central1, subject to change.*

Google Cloud

Jellyfish Pictures

Award-winning visual special effects (VFX) and animation studio

CUSTOMER STORY

JELLYFISH PICTURES

> About

Jellyfish Pictures creates visual effects (VFX) and animation for advertising, film, and television. The company has won Emmy, VES, and BAFTA awards.

> Challenges

Jellyfish planned two new animation studios in London. The company didn't want to pay London real-estate prices to house IT infrastructure. Another dilemma: how to provide desktops to freelancers who work outside the studio.

> Solution

Jellyfish built a private cloud in an existing location to host virtual desktops. Artists in the new studios log in to their virtual desktops from Dell Wyse Zero Clients, powered by Teradici PCoIP technology. Freelancers use Teradici Cloud Access Software to access virtual desktops hosted in Microsoft Azure.

> Results

Avoiding on-site IT infrastructure saved ~$100,000. Now that Jellyfish can quickly provision virtual desktops for freelancers in any location, the company can take on larger projects with tighter deadlines. Creative content remains secure because it never leaves the cloud.

VFX Studio Reference Architecture

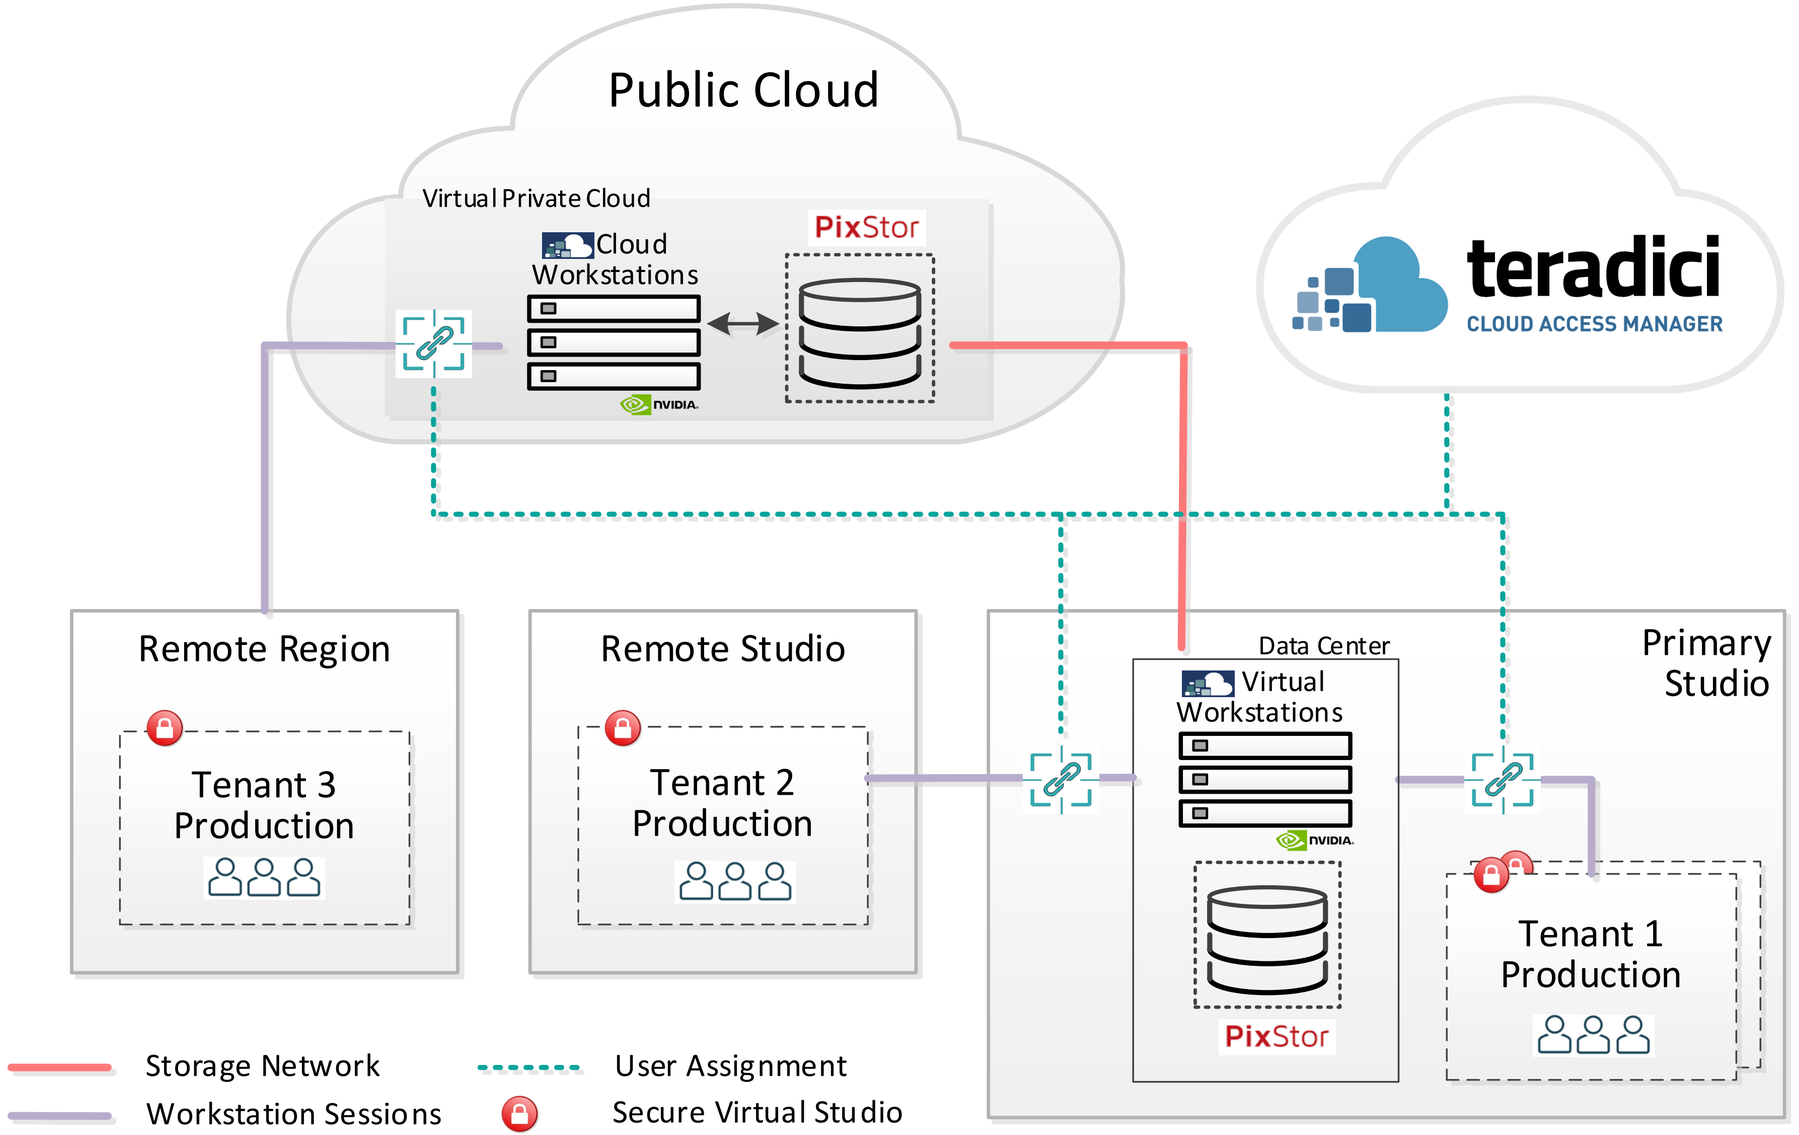

Tenant Isolation using VLANs and Security Zones

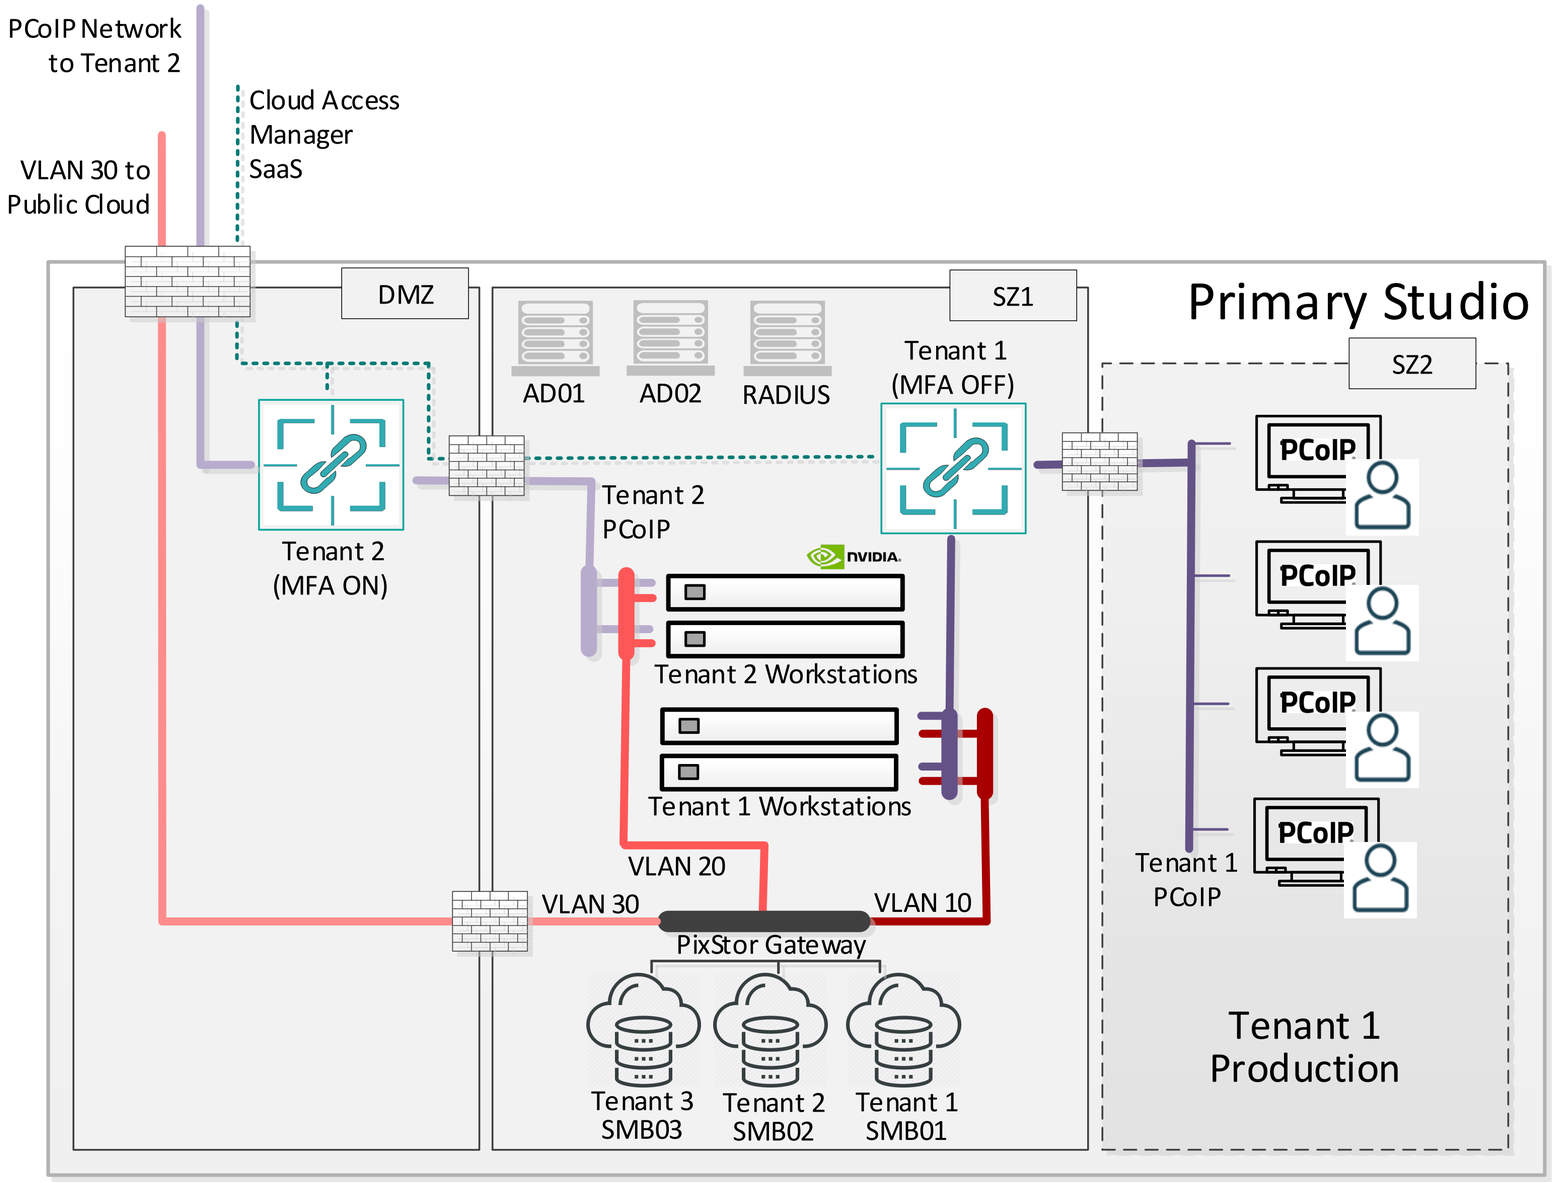

Storage Isolation using SMB Containers

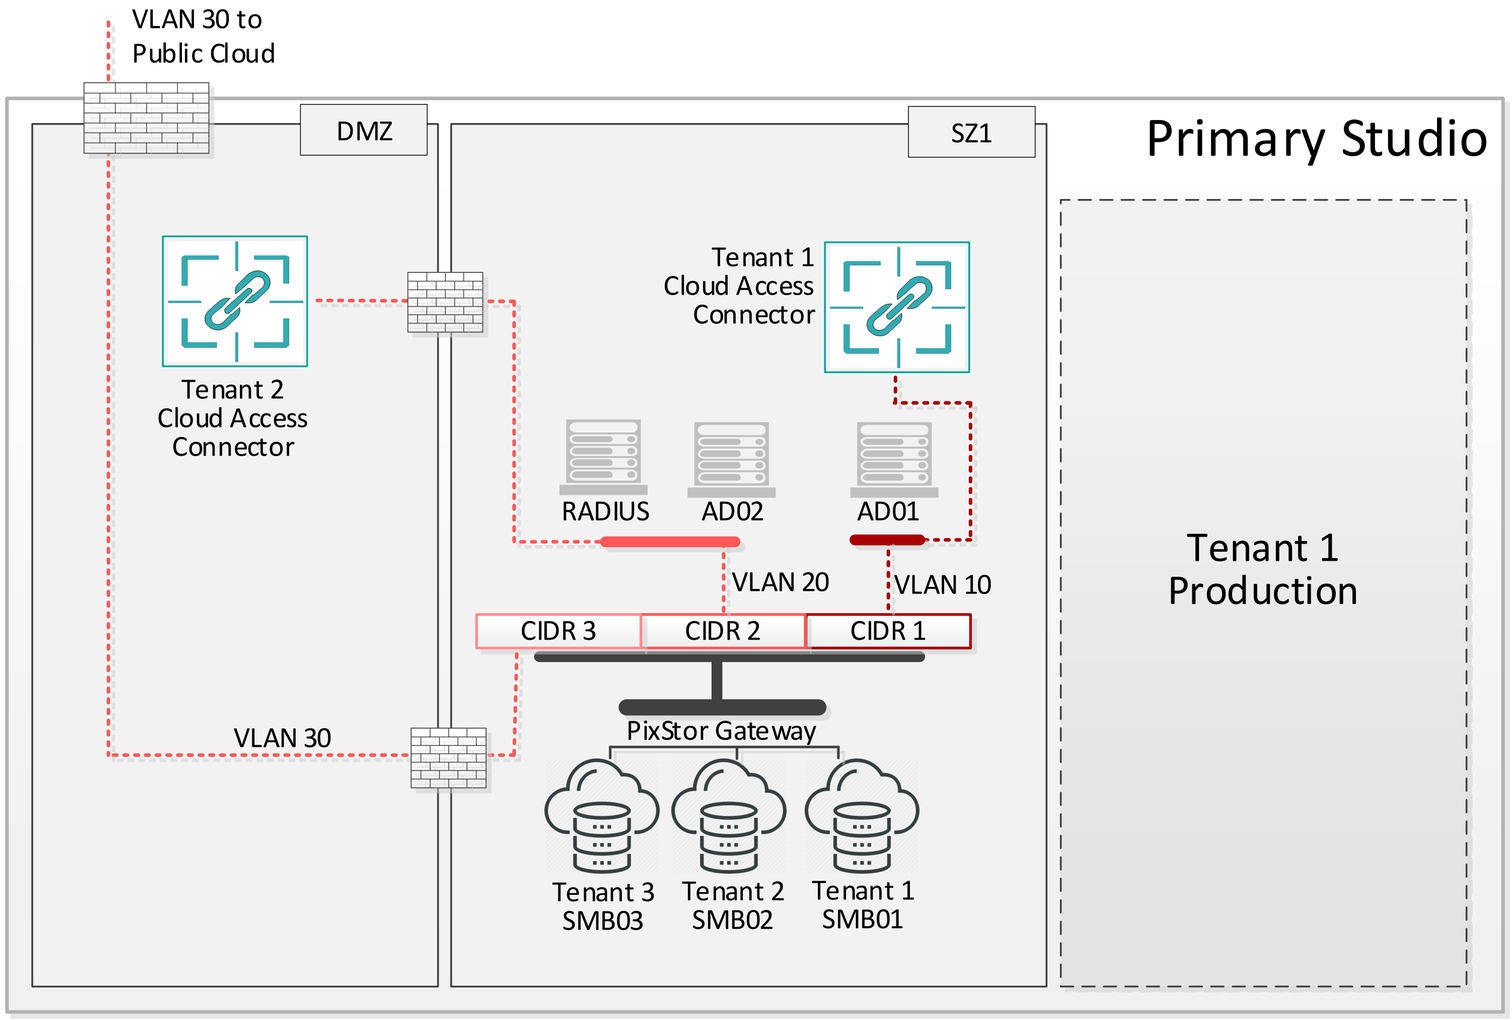

PixStor Gateway Services:

SMB Container 01
IP = CIDR Block 1
VLAN TAG = 10
Director Service = AD01
Mount = mmfs/tenant1

SMB Container 02
IP = CIDR Block 2
VLAN TAG = 20
Director Service = AD02
Mount = mmfs/tenant2

SMB Container 03
IP = CIDR Block 3
VLAN TAG = 30
Director Service = AD03 (Cloud)
Mount = mmfs/tenant3

Production Network Isolation from Public Internet

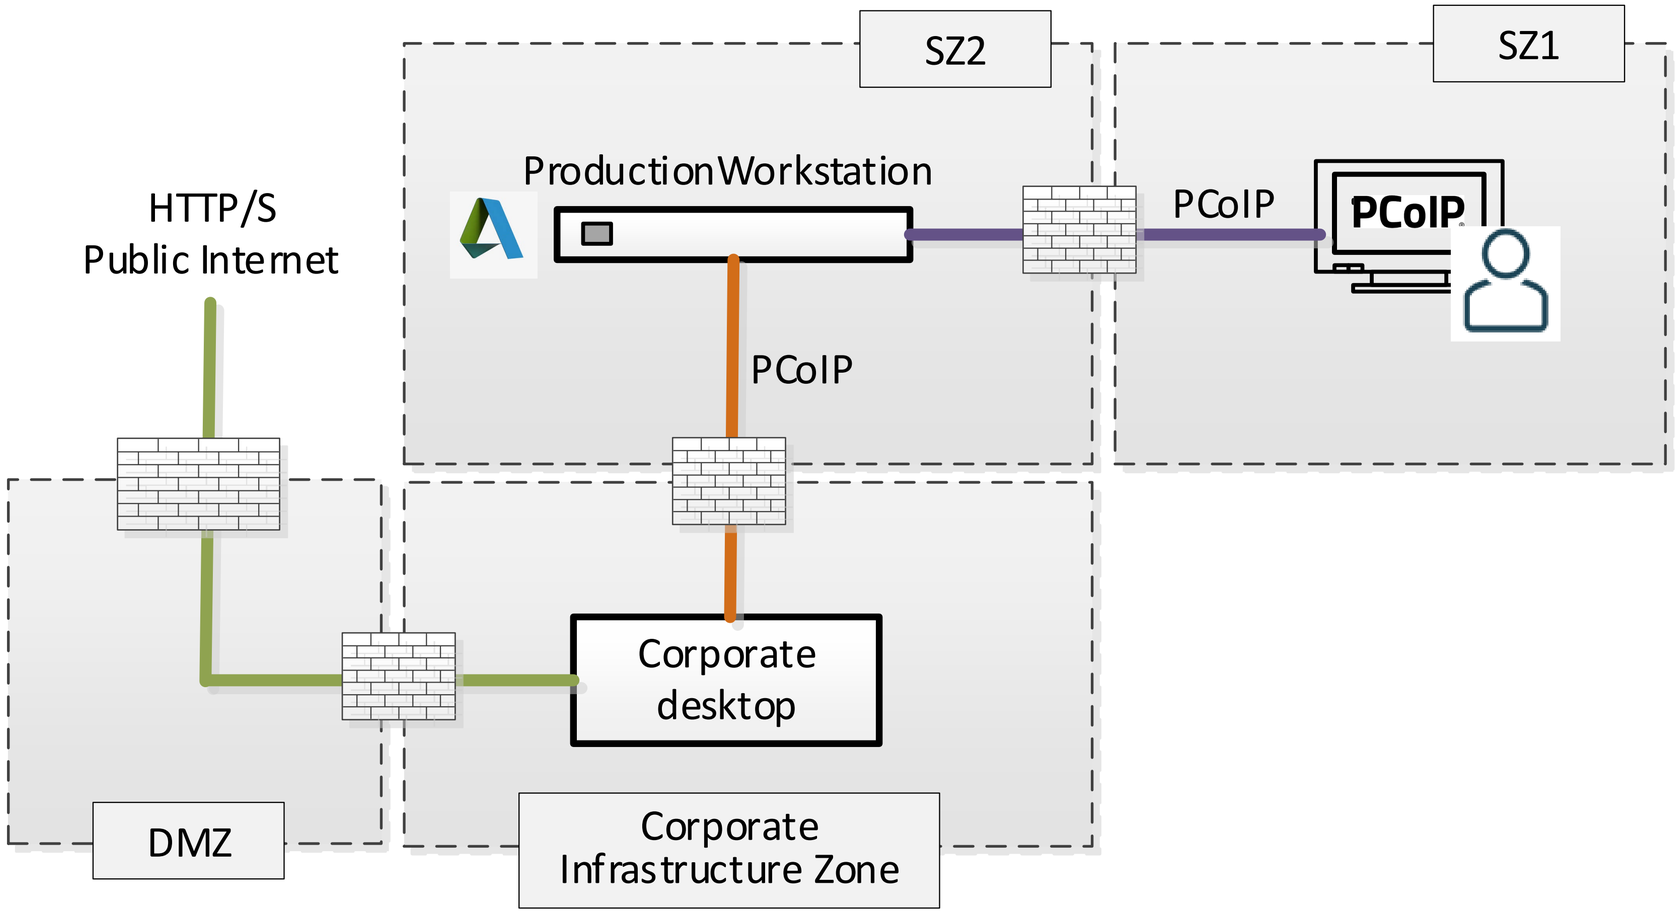

DNEG

Visual effects leader

CUSTOMER STORY

DNEG

About

Headquartered in London, DNEG is one of the world's leading visual effects (VFX), animation, and stereo conversion companies for feature film and television. DNEG has nine offices in four countries.

Challenges

Preparing new studio space was costly and time-consuming because of artist workstation power and cooling demands. The powerful workstations were also costly—and noisy.

Solution

DNEG built an on-premises virtual workstation infrastructure. Artists access their virtual workstations from PCoIP Zero Clients using Teradici Cloud Access Software. Compared to other solutions DNEG evaluated, Teradici Cloud Access Software provided better accuracy, compression, and support for Wacom (pen-interface) tablets

Results

Moving to the permanent Montreal studio was fast and simple: no special power supplies, no expensive chillers. Opening smaller studios more quickly and at less cost gives DNEG a new competitive advantage. TCO dropped because moving to virtual workstations lets DNEG oversubscribe CPU and GPU resources. And artists like the solution because the office is quieter and they didn't have to change the way they worked.

Remote Session Planning – Network Bandwidth Considerations

Network Bandwidth Demand:

- A variety of contributors to network bandwidth
  - Display topology & resolution (e.g. single vs. dual monitor; HD vs. UHD resolution)
  - Content & protocol frame rates (e.g. 30 fps vs. 60 fps)
  - Workload and artist role (content creation vs. content playback)
  - Protocol quality settings
- We recommend starting with a Proof of Concept (PoC) using a representative set of users

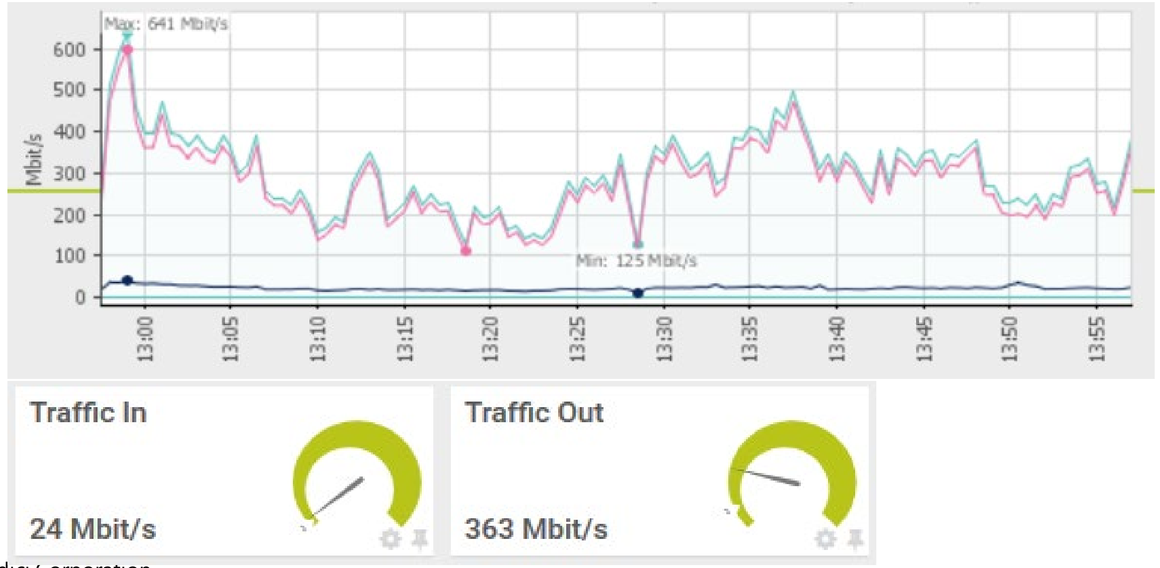

Studio Example Configuration:

- 69 x Dual HD displays
- Selection of VFX apps including 3ds Max, Maya, etc.
- 363 Mbps accumulated average bandwidth:
  - 5.3 Mbps per user on average

Remote Session Planning – Network Latency Considerations

Latency – Distance from Datacenter:

- High interactivity is very important for artists using Wacom tablets
  - **If network latency < 25ms**: Great user experience when tablet USB is directly ‘bridged’ to host tablet drivers via PCoIP
  - **If network latency > 25ms**: User experience is improved if tablet USB is ‘locally terminated’ at client
  - ‘Local termination’ presents an immediate local cursor to assist hand-eye coordination in addition to bridging parameters such as pen position, pressure and angle

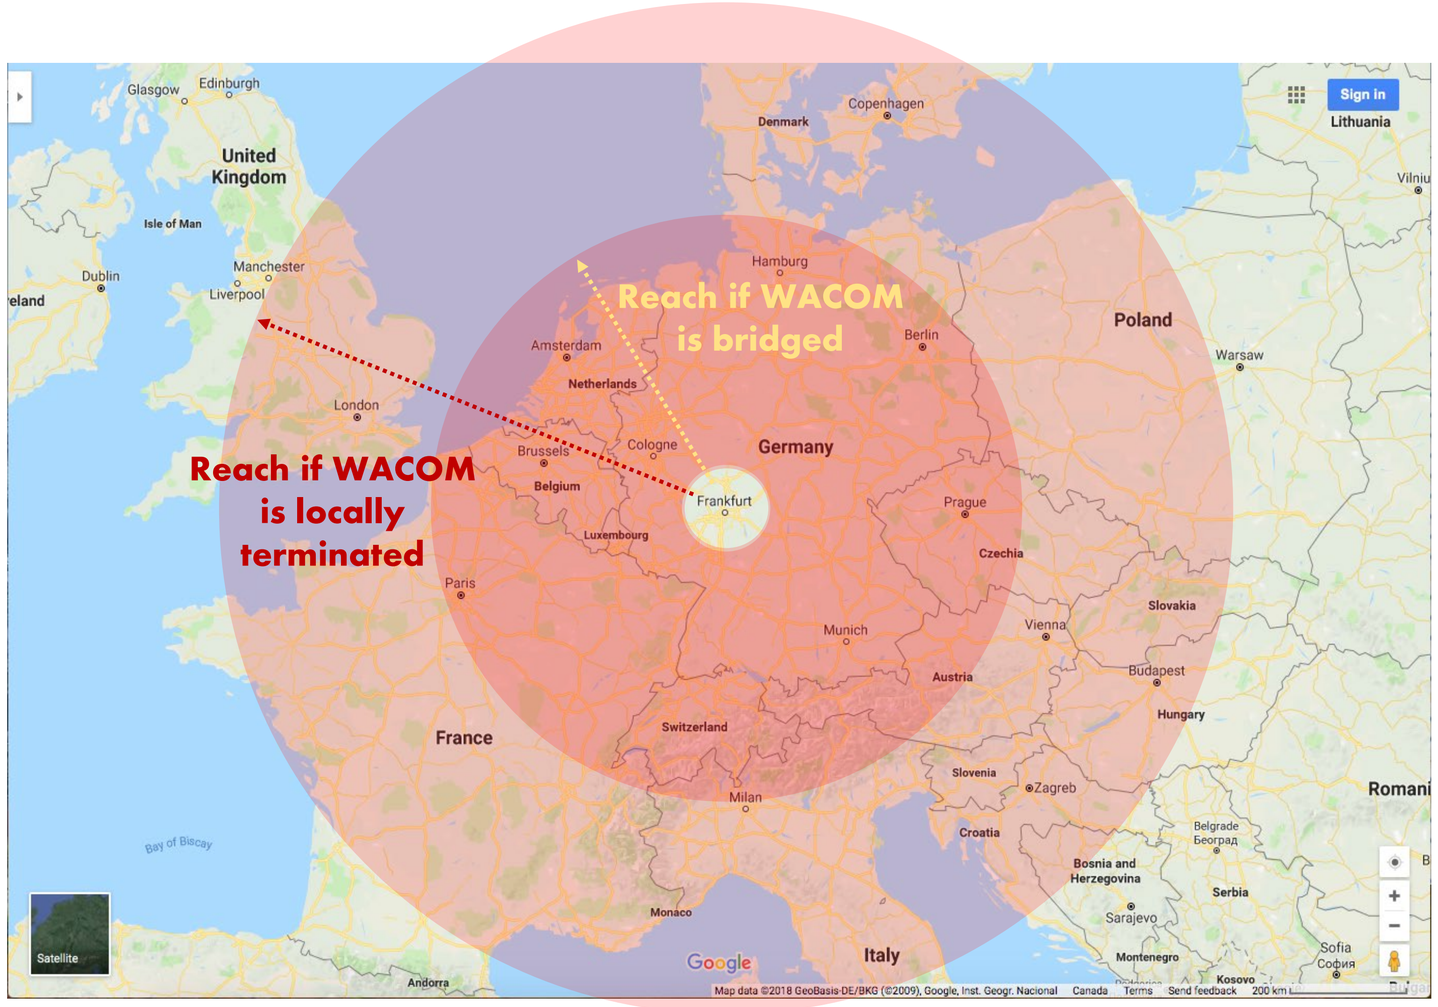

THANK YOU!