Running 3D Applications on Remote HPC Infrastructures

Paolo Maggi <paolo.maggi@nice-software.com>
R&D Manager

About NICE

- **Pioneer in Technical and Engineering Cloud solutions**
  - Focused on enterprise Grid/Cloud solutions throughout all industries since 1996
  - Worked closely with ISV and HW vendors since the beginning
  - Strong relationship with Research and Academia
- **Core business: Access to Grid / HPC / Cloud solutions and Remote Visualization**
  - NICE EnginFrame Technical Computing Portal
  - NICE DCV Remote Visualization technology
- **Other relevant competencies**
  - Distributed Resource Management
  - Grid Intelligence
  - Big Data

NICE

NICE Customers and Market Segments

Energy & Chemical

Anadarko, BGP, British Gas, Cairn Energy, ConocoPhillips, Daqing Oil Field, Dow Chemical, Dupont, ENI/AGIP, EOG Resources, FMC, GEA Process Eng., GXT, Huabei Oil Field, Maersk Oil, Nova Chemical, Pentair, PetroChina, Tuha Oil Field, Tullow Oil, Xinjiang Oil Field

Life Sciences and Medical

Baxter, Eli Lilly, Howard Hughes Medical Institute, Institut Pasteur, Novartis, Ontario Cancer Institute, Shanghai Inst. for Biological Sciences

Research & Education

CINECA, DLTM, Insis, NCSA, South Dakota State University, TU Ilmenau, TU Ostrava, Uni Birmingham, Uni Buffalo, Uni LETI, Uni Liverpool, Uni Nagoya, Uni Padova

Aerospace & Defense

1st Aerospace Institute, Alenia Aermacchi, AVIC, BEMEI, Boeing, DLR German Aerospace, Procter&Gamble, P&W Canada, SelexGalileo, Gulfstream, MBDA, NORDAM, NGC, Parker Hannifin, Sikorsky

Automotive & Manufacturing

3M, ABB, Audi, BMW, Bosch, Continental, Daimler, Delphi, Emerson, Global Castings, Hyundai, JLR, Knowbe, Liebherr, Magneti Marelli, Mazda, McLaren, Panalpina, Rolls Royce, Scania, Toyota, Volkswagen

Others

Ansys, AWE, BlueCloud, CADFEM, CFX Berlin, Silicon Labs, Samsung, Tycoelectronics, Westinghouse

Key Partners

NVIDIA
sgi
hp
IBM
DELL
FUJITSU
lenovo
Platform Computing an IBM Company
Adaptive Computing
UNIVA
Schlumberger
ANSYS
MSC Software
esri
vsg Visualization Sciences Group
CD-adapco

Running 3D applications on remote HPC infrastructures

Why?

A Typical Scenario in the CAE sector

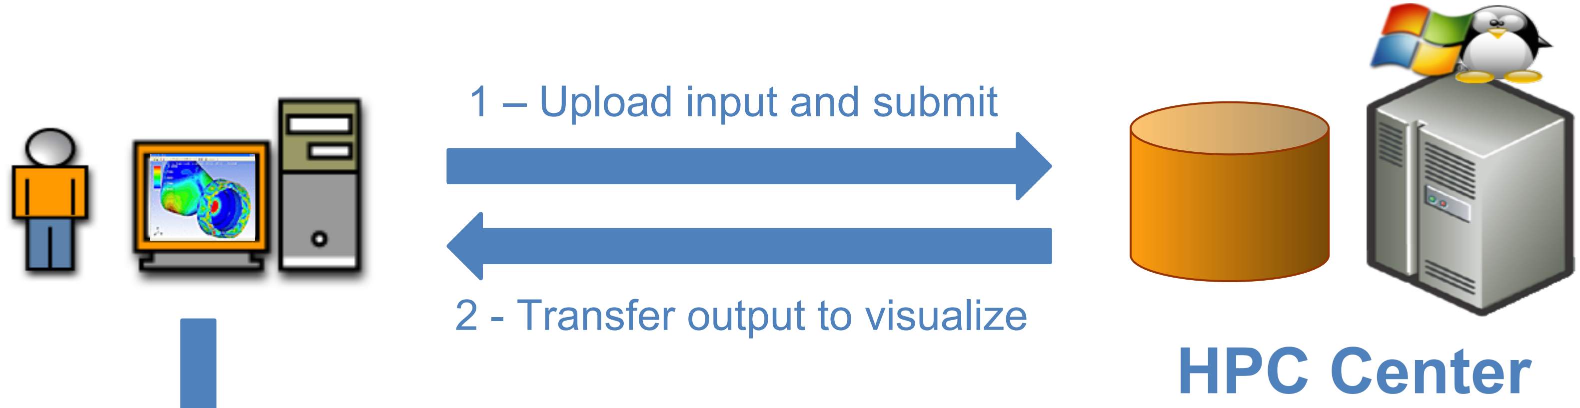

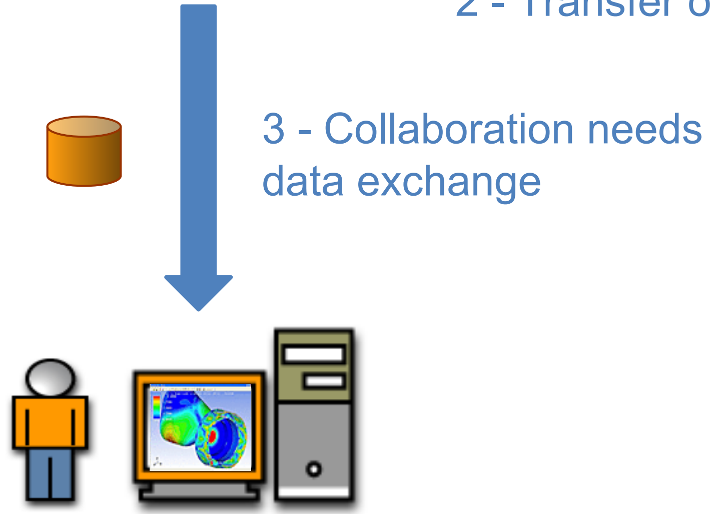

Every engineer has one or more **engineering workstations** to:

- Run pre/post processing tools (high-end GPU required for 3D graphic rendering)
- Run small to medium serial analysis

Big companies have tens to hundreds of engineers

A Typical Scenario in the CAD sector

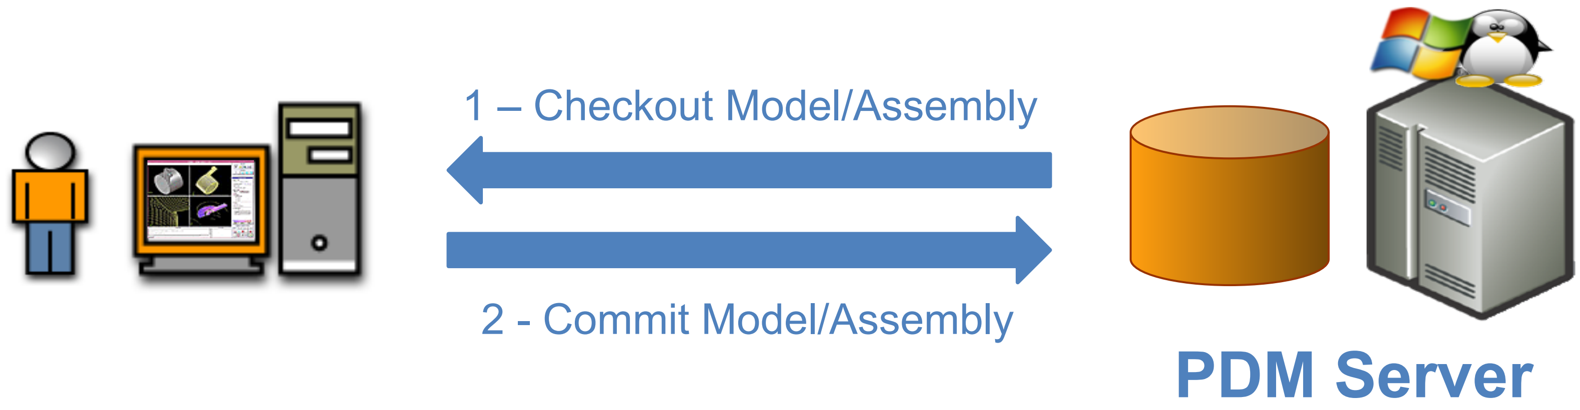

Every designer has its own **CAD workstations** to:

- Run CAD software (high-end GPU required for 3D graphic rendering)
- Run analysis tools (e.g. clash)

Big companies have tens to thousands of engineers

Issues with the "workstation model"

I love my workstation… what is wrong with it?

NETWORK
Network overload leading to poor performance and response times all round

COST
Expensive, dedicated workstations (GPU, memory, …) with short lifecycle

IT MANAGEMENT
Support, update and replace tens to tens of thousands of workstations

WORKSTATION SIZING
Workstations have to be sized for the largest expected models

SECURITY
Moving sensitive data around (in/out organization) is always risky

WORKFORCE
Current models do not support a diverse mobile workforce

Let’s Think Different

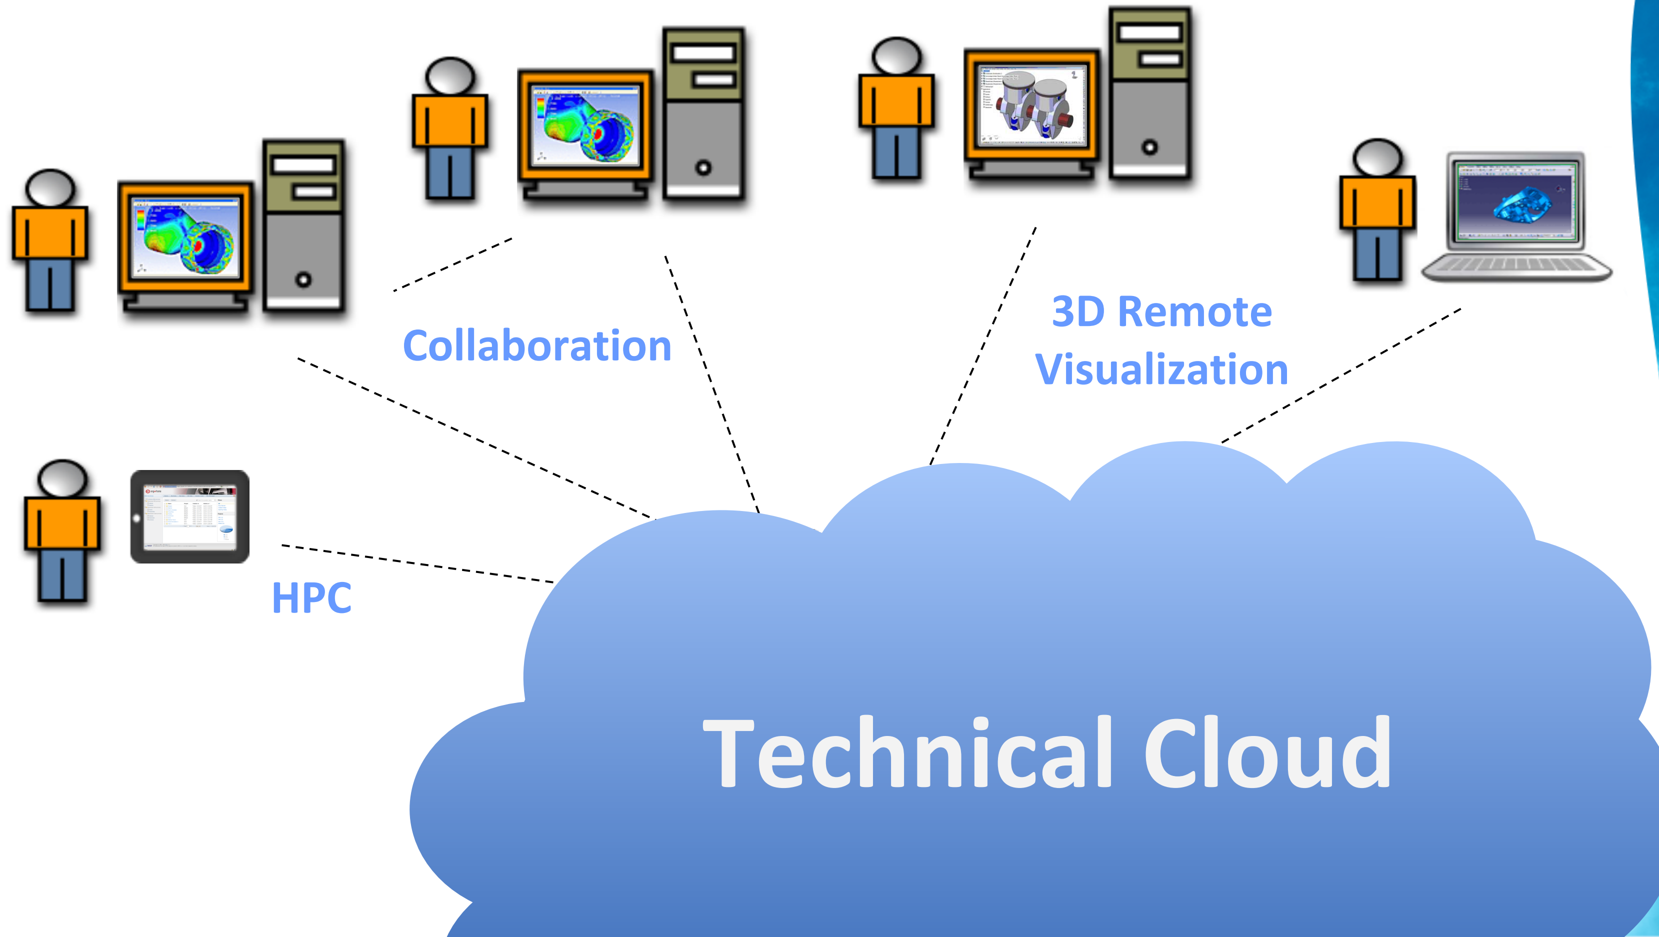

Technical Clouds

What is a Technical Cloud?

A technical cloud enables convenient, on-demand network access to a shared pool of computing resources that run:

- **engineering simulation tools** and other **HPC applications**
- **high-end 3D technical applications** (like visualization applications for scientific data, CAD applications, etc.)

Benefits of the Technical Cloud

… but how a Technical Cloud can help me?

NETWORK
Network is no more a bottleneck and data loads faster

COST
Centralized & Shared servers are less expensive to buy & manage

IT MANAGEMENT
Support, update and replacement are more efficient & do not affect users

WORKSTATION SIZING
Resources are dynamically sized based on users needs

SECURITY
Sensitive data remain within protected data center with full access control

WORKFORCE
Users can virtually connect & collaborate from anywhere with any client

Enabling Technologies

- Technologies providing secure remote network access to interactive 3D applications leveraging server-side graphic hardware acceleration (GPUs)

- A software stack that allows the end-user to easily launch and access remote interactive applications and takes care of managing and load balancing applications and desktop sessions running within a Visualization Farm

NICE Proposal for Technical Clouds

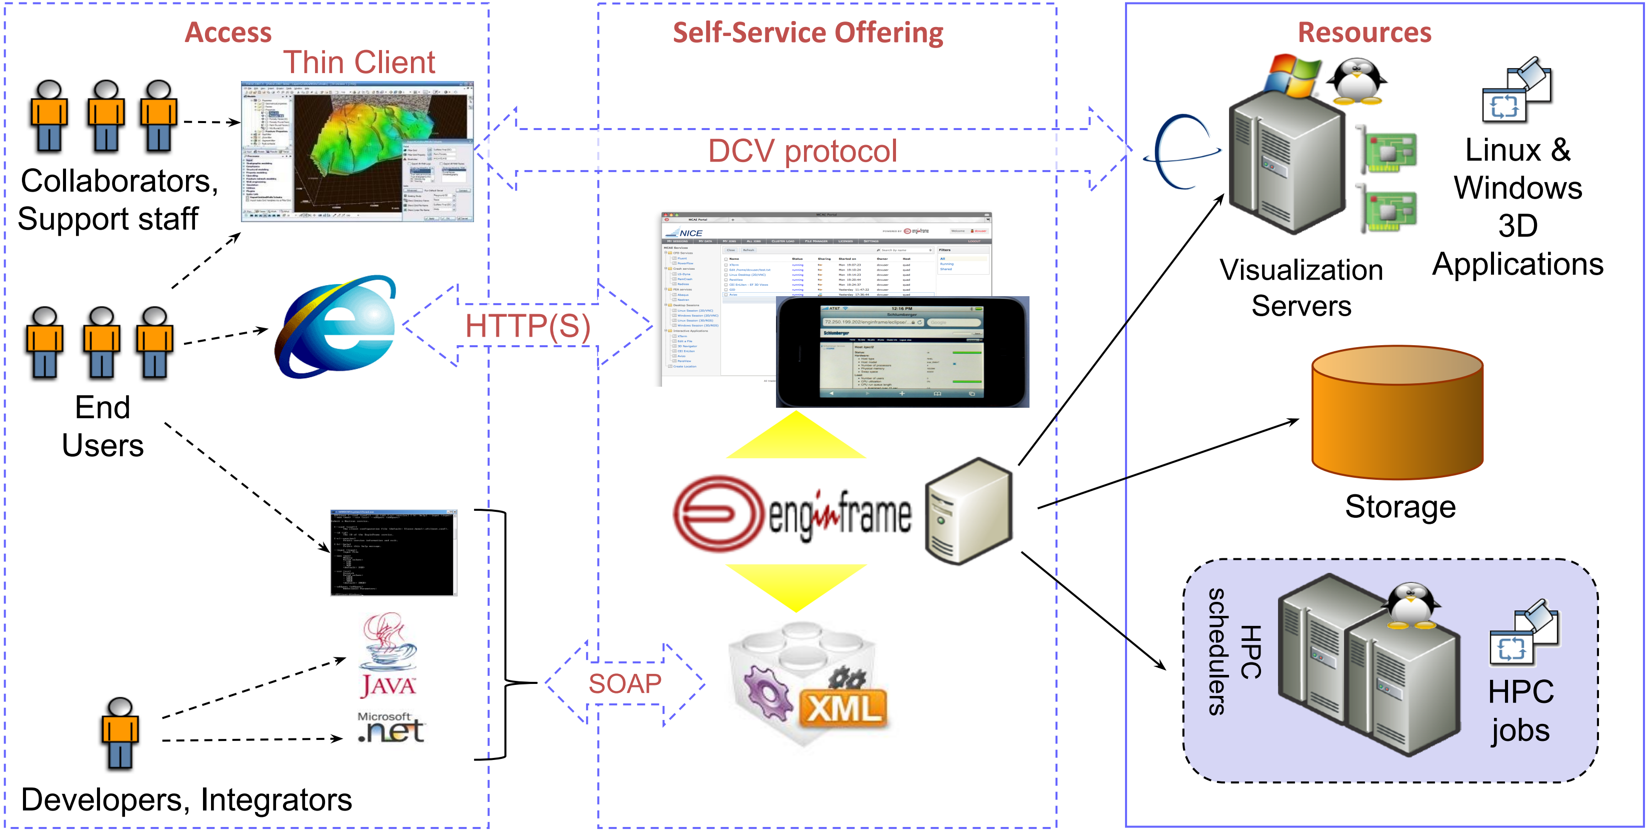

NICE EnginFrame – HPC and 3D Cloud Portal

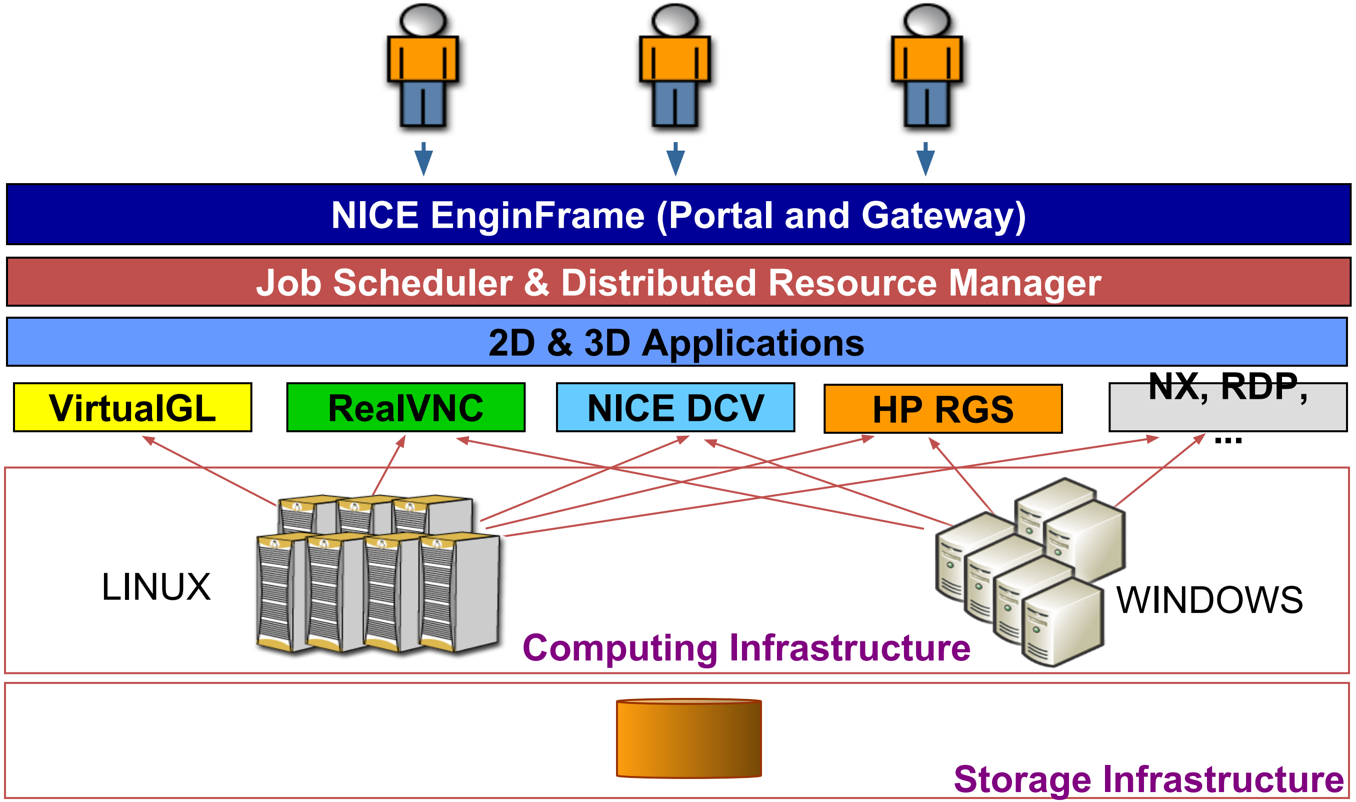

How Does It Work?

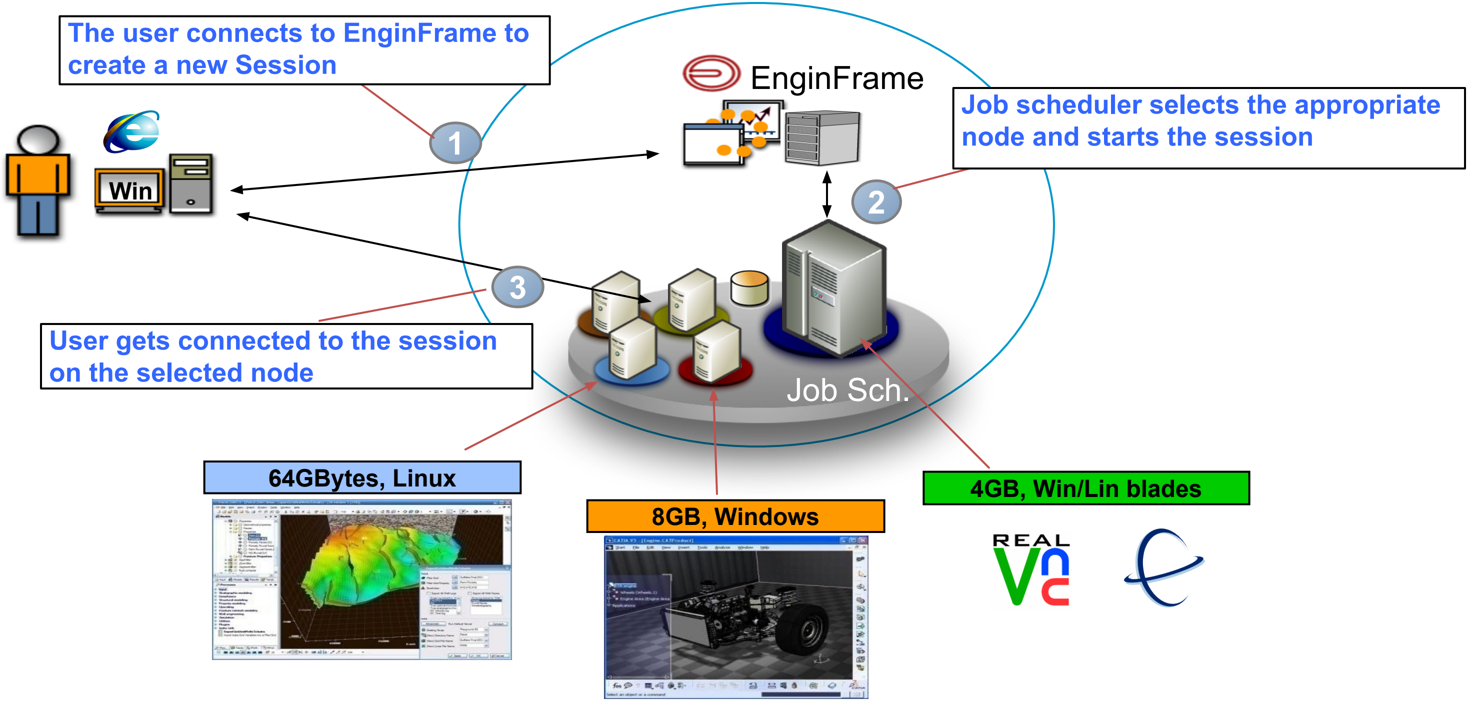

Heterogeneous infrastructure: HW, OS, middle-wares

Benefits for the End Users

- **Access to Applications as a Service**
  - Hide infrastructure details (site, platform, etc…)
  - Session allocation can be influenced by memory requirements, data location affinity and other customer-specific parameters
  - Intelligent load balancing of sessions, based on the Job Scheduler
    - Memory-aware, Memory reservation, Application license-aware
  - Fewer or better optimized data transfers
- **Collaboration, session sharing**
  - The owner can invite a colleague to join a session, without disclosing the user's password
- **Easy management of active session**
  - Create new sessions with user-specific preferences (resolution, etc...)
  - List, reconnect and kill existing sessions
- **Seamless Access to the sessions and Single Sign-On**
  - Automation of session-level password creation and destroy (e.g. login via NTLM / ActiveDirectory credentials and map to Linux user)

Benefits for the Administrators

- **Increased level of service provided to users**
  - Sessions are load balanced by the Job Scheduler to match user needs
  - Memory reservation and Data locality scheduling
  - Reduce help desk calls
  - Exposed services can be personalized per user/group/project
- **Accounting**
  - Sessions are jobs, so the resource usage accounting by user, group, project can be collected through any Analytics tool
- **Monitoring**
  - The load and usage of the login farm is monitored via EnginFrame
  - Node loading conditions, active sessions
  - Administrators can control and manage users' idle or stuck sessions
- **Support**
  - Administrators can connect to user's sessions to provide support
- **Security**
  - Easy integration into identity services, SSO, Enterprise portals

About NICE Desktop Cloud Visualization

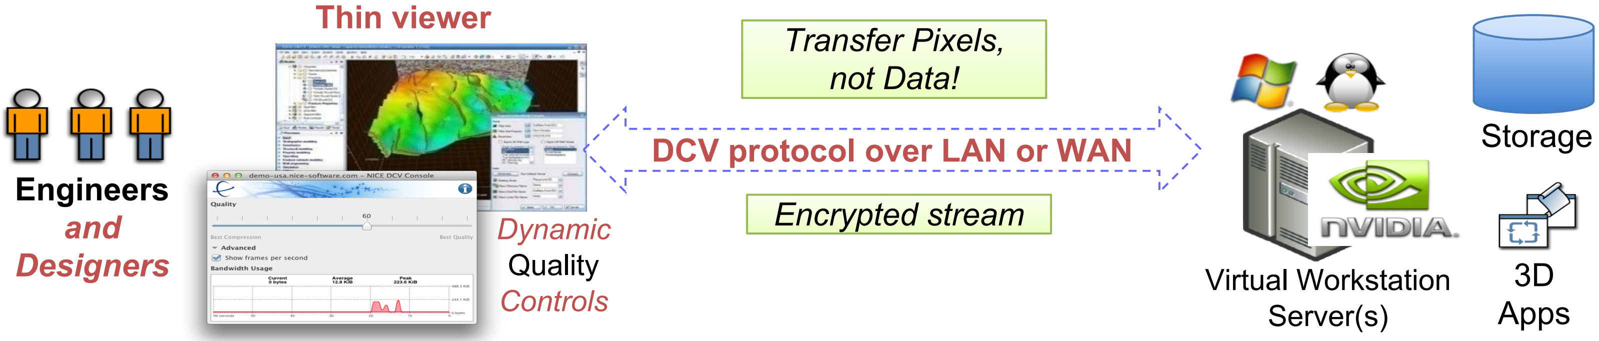

- Provides **efficient remote access** to graphic-intensive, professional OpenGL and DirectX applications
- A **central server** is accessed by remote users:
  - requires an high-end **NVIDIA GPU** (Quadro and Grid families) on the server side
  - the server host can be a **physical** or **virtual machine** (KVM, Xenserver, ESXi)
- Users only need low-end machines with network connectivity to interact with remote application
- One or more users (collaborators) can simultaneously access the server
- Supports a broad range of OSes and devices:
  - **Server**: Windows 7, XP, 2008, RHEL, SLES, CentOS, Oracle Linux
  - **Client**: Windows, Linux, OS X, iOS, Android, applet
- Allows to **share GPUs** among multiple users and OS.
- Provides **dynamic bandwidth optimization** balancing quality Vs. frame rate
- Validated and optimized for Technical Computing applications

Main Protocol Characteristics

- **Only final rendered images (not geometry and scene information) are transmitted to the client**
  - Provides insulation and protection of proprietary customer information
- **Automatically adapts to heterogeneous networking infrastructures like LAN, WAN and VPN**
  - Automatically deals with bandwidth and latency constraints
- **Users can specify the compression level used to send the final image to the endstation**
  - Balances quality vs. frame rate on low-bandwidth conditions
- **High network latency tolerance**
  - Good results with RTT up to 120-150ms
- **Supports both <u>lossless</u> and <u>lossy</u> compression**
  - Lossless compression is very important for some markets like the medical imaging sector
  - Since version 2014.0 uses **H.264** codec.
- **Support for high quality updates when network & processor conditions permit**
  - Still images receive a “quality boost” so that they are always “pixel perfect”
- **Transmitted data can be encrypted using 128/256-bit AES cypher**
- **Supports HTTP and SOCKSv5 proxies**

H.264-based Encoding

- NICE DCV compresses and encodes frames before sending them to the clients
- NICE DCV 2013.0 uses an JPEG-based encoder (still images codec)
- NICE DCV 2014.0 will use a H.264-based encoder
  - a fallback mechanism to the old codec is used with GPUs/clients that do not support H.264 encoding
- H.264 is a video codec
  - it can reduce the bandwidth utilization (roughly speaking) by sending only the differences between a frame and the previous one
- H.264 can reduce the bandwidth utilization by a factor of 5 or more (the actual reduction depends on the specific images that are being transmitted)

H.264-based Encoding Bandwidth Utilization

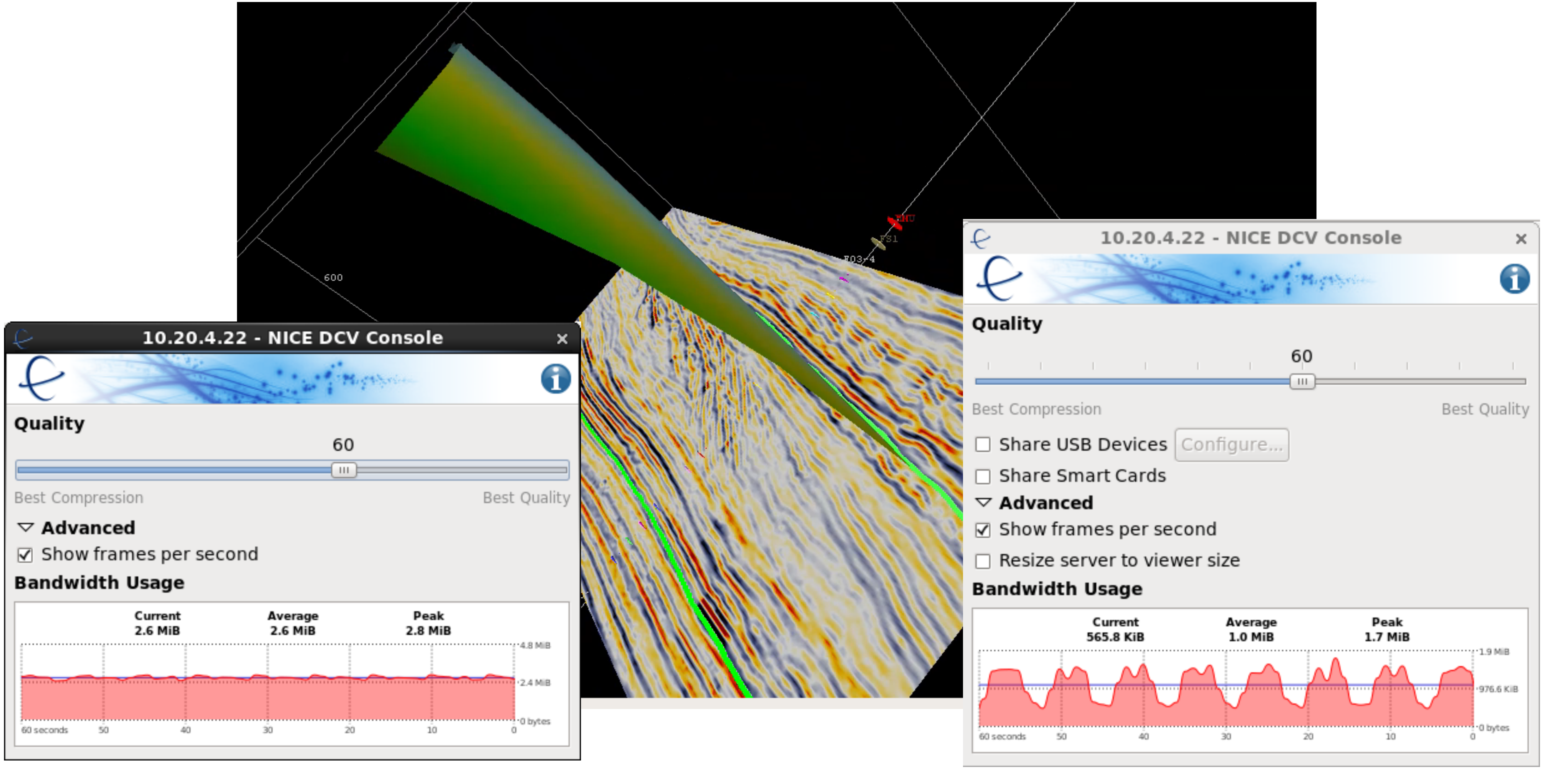

DCV 2013.0

DCV 2014.0

NVIDIA GRID SDK

- The NVIDIA GRID Software Development Kit is a comprehensive suite of tools for NVIDIA GPUs that enables high performance graphics capture and the usage of the NVIDIA H.264 hardware encoder
- NICE DCV uses NVIDIA GRID SDK to:
  - efficiently capture frames (reduces latency due to frame buffer reads)
  - use the NVIDIA H.264 hardware encoder (reduces latency due to frame encoding)

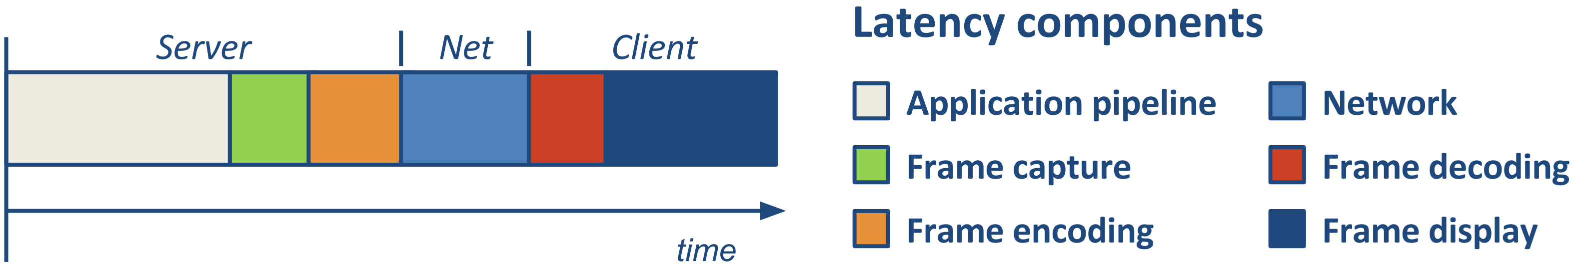

NVIDIA H.264 Hardware Encoder - NVENC

- NVIDIA's latest generation of GPUs, based on the Kepler architecture, features a hardware-based H.264 video encoder.
- The hardware is optimized to provide excellent quality at high performance
  - up to 120 frames per second of 1920 × 1080 progressive video in high-quality (HQ preset) mode
- NVENC, being dedicated H.264 hardware on the GPU chip, does not take away processing power from other tasks that can be performed on the CPU and GPU's graphics engine
- Using NVENC, NICE DCV 2014.0 has better performances and consumes less CPU power
  - In the previous versions, frame encoding was performed on the CPU

Multiple Linux Sessions with GPU Sharing

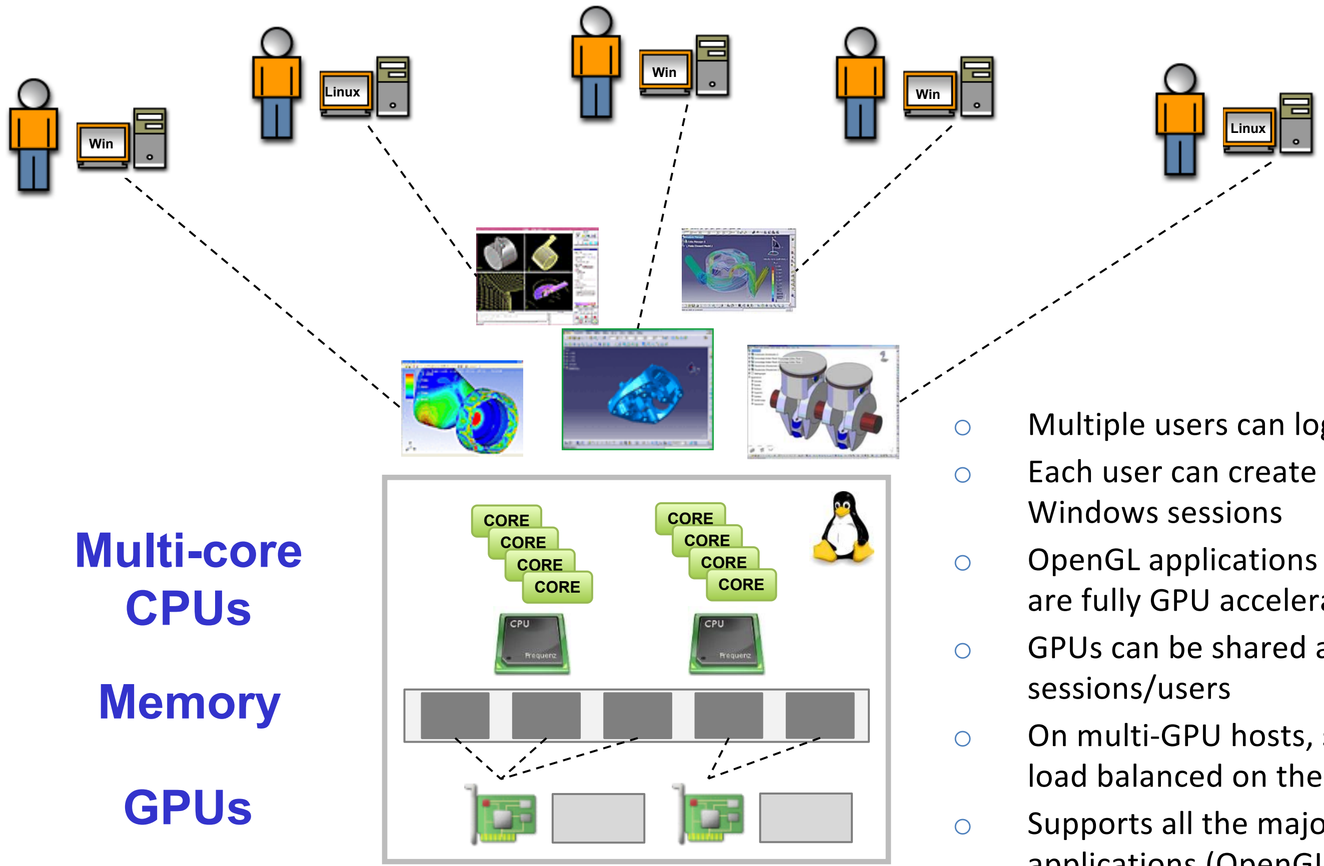

- Multiple users can log into a single host
- Each user can create one or more (isolated) X Windows sessions
- OpenGL applications running in the sessions are fully GPU accelerated
- GPUs can be shared among multiple sessions/users
- On multi-GPU hosts, sessions are automatically load balanced on the available GPUs
- Supports all the major OpenGL technical applications (OpenGL 3.x-4.x support)

High Level Architecture (Linux)

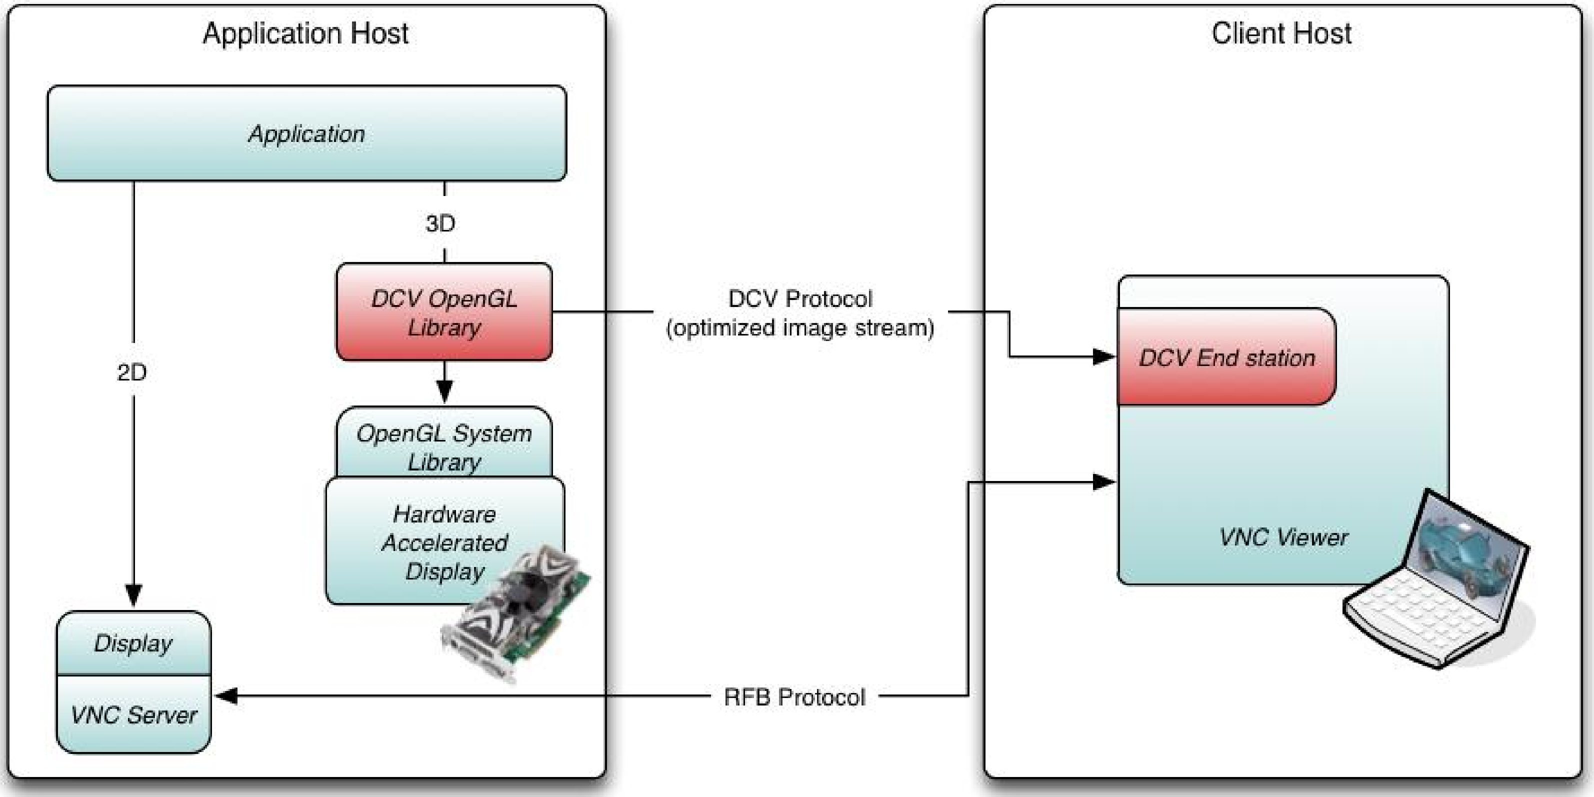

OpenGL/DirectX-based Windows Applications

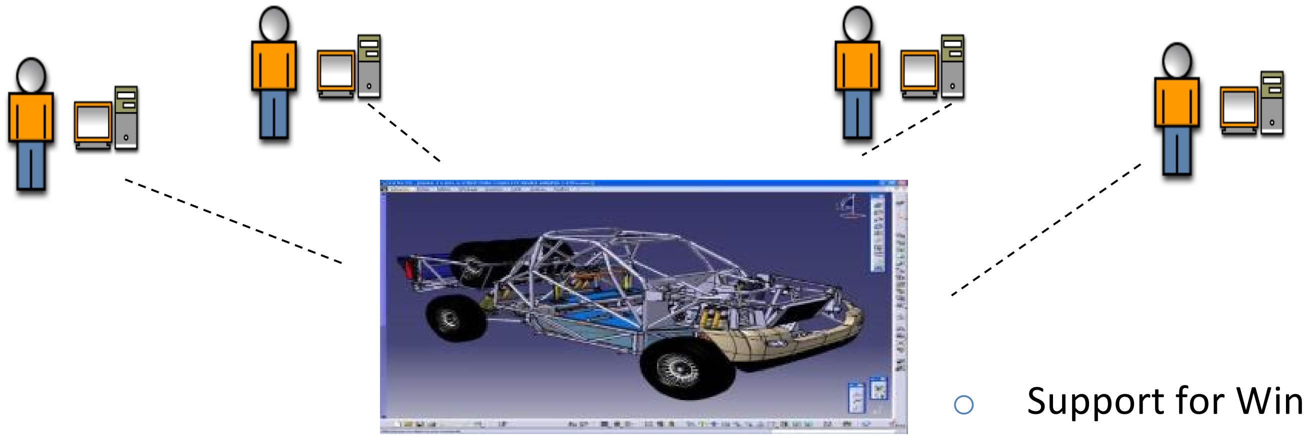

Multi-core CPUs

Memory

GPU

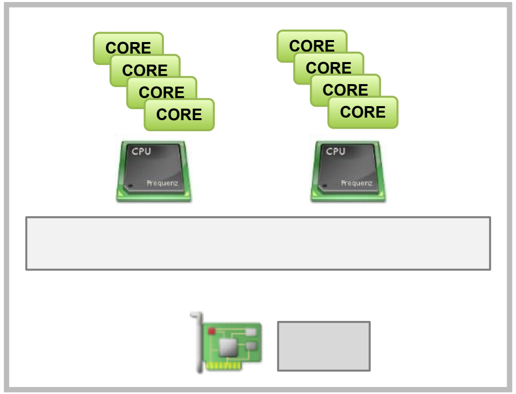

- Support for Windows physical hosts
  - One session per host
- Supports all the major technical applications:
  - OpenGL 1.x, 2.x, 3.x and 4.x
  - DirectX 9, 10 and 11
- One or more users (collaborators) can simultaneously access the session

Virtual Machine Remotization

- NICE DCV can remotize OpenGL/DirectX applications running inside a VM

- Supports different technologies:
  - GPU pass-through: XenServer, ESX, KVM
  - vGPU: Citrix XenServer 6.2+ and VMware vSphere ESXi 6.0 (requires NVIDIA GRID cards)
  - NICE «External Rendering Server» technology: KVM, XenServer, ESX

First product on the market to allow the sharing of physical GPUs between multiple Windows VMs while maintaining full OpenGL acceleration and workstation-class performance.

Allows to share GPU among multiple users and multiple Operating Systems.

NICE DCV with GPU pass-through

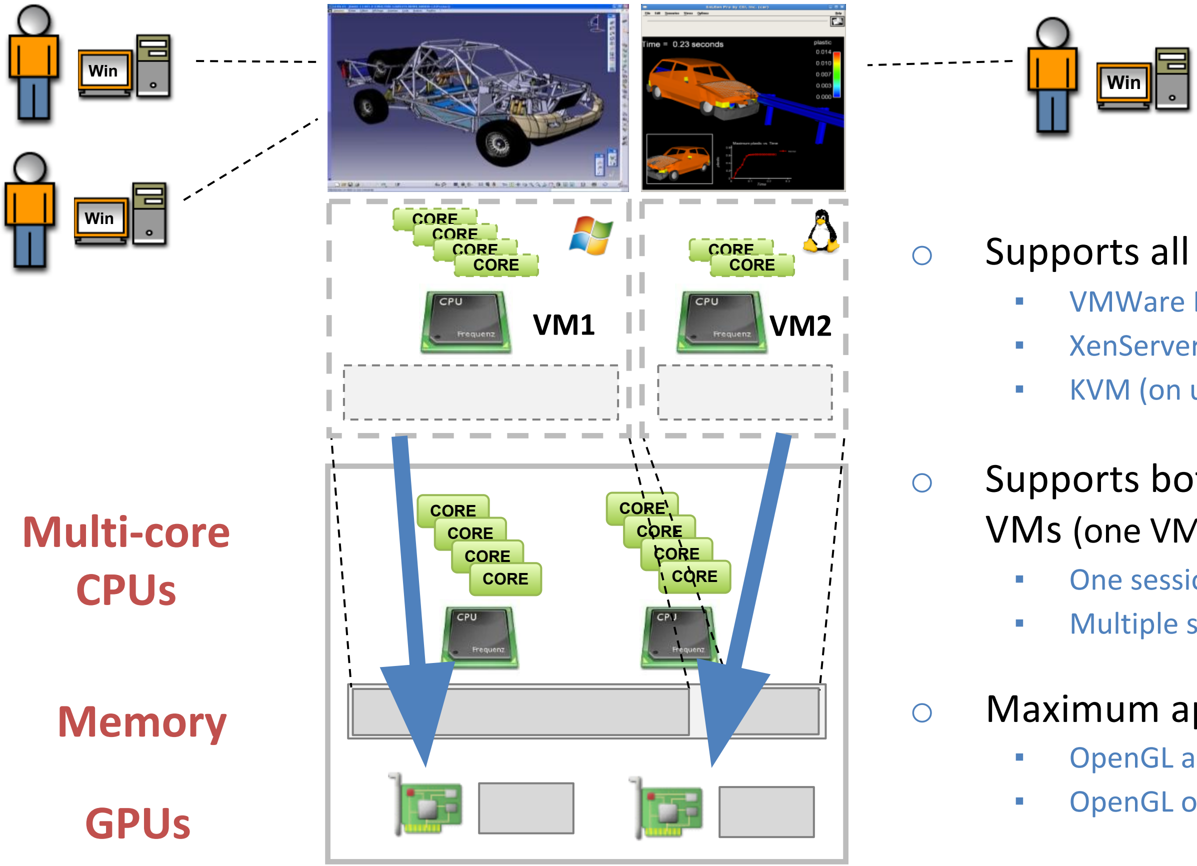

- Supports all major hypervisors:
  - VMWare ESX
  - XenServer
  - KVM (on upcoming RHEL 7)
- Supports both Windows and Linux VMs (one VM per GPU)
  - One session per VM on Windows
  - Multiple sessions per VM on Linux
- Maximum application compatibility
  - OpenGL and DIrectX on Windows
  - OpenGL on Windows

NVIDIA vGPU Technology

- NVIDIA GRID vGPU enables multiple virtual machines (VMs) to have simultaneous, direct access to a single physical GPU:
  - high graphic performances
  - high application compatibility

- Limitations:
  - Works only with NVIDIA GRID GPUs (e.g. NVIDIA K2)
  - Today, the only supported hypervisors are XenServer 6.2 SP1 or above and VMware vSphere ESXi 6.0
  - Today, Linux is not supported
  - Applications using CUDA or OpenCL are not supported

- NICE DCV 2014.0 fully supports NVIDIA vGPU technology

NICE DCV with vGPU

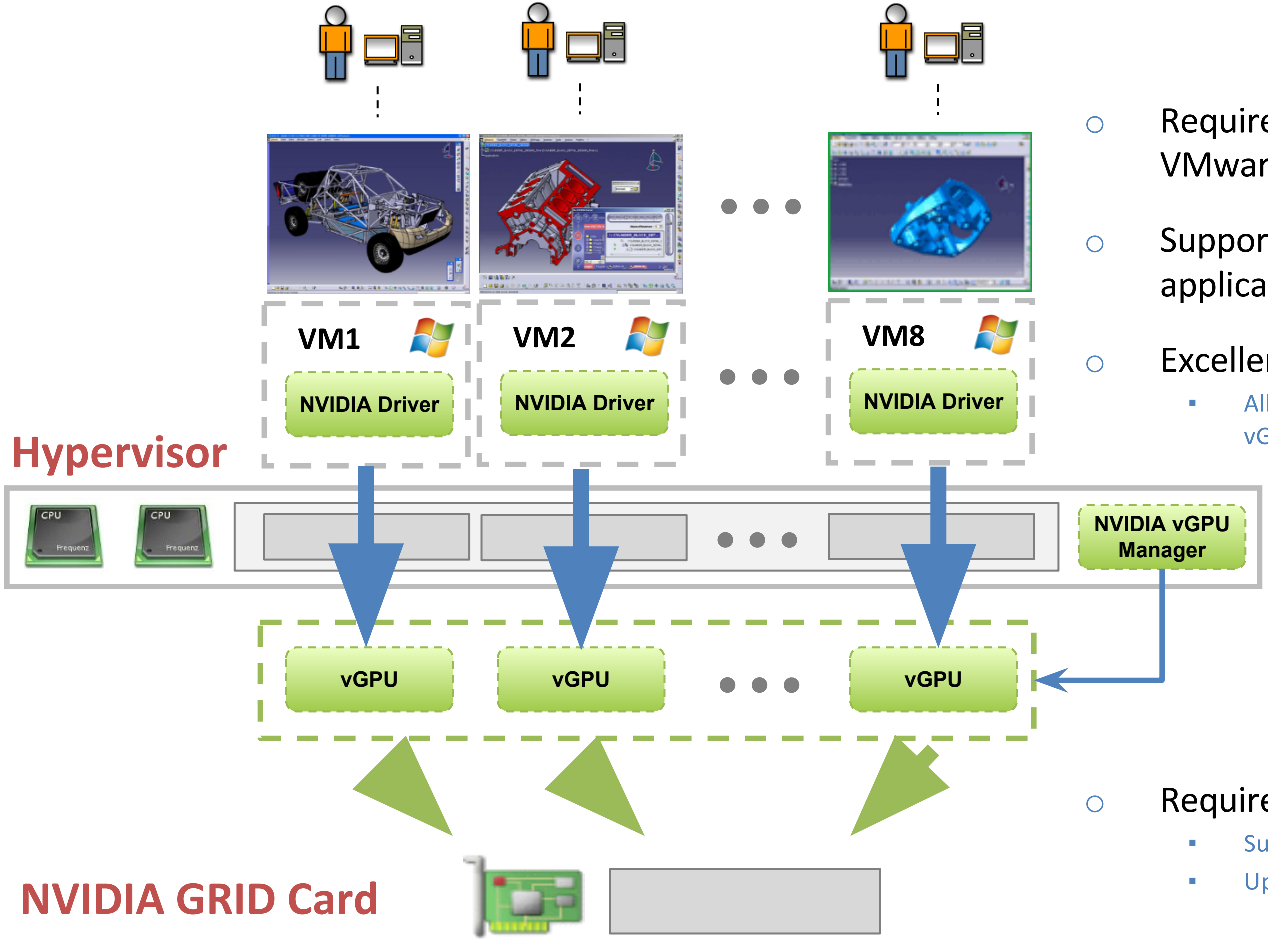

- Requires Citrix XenServer 6.2+ or VMware vSphere ESXi 6.0
- Supports OpenGL and DirectX applications on Windows 7
- Excellent application compatibility
  - All applications running with NVIDIA vGPU
- Requires an NVIDIA GRID Card
  - Suggested NVIDIA K2
  - Up to 8 VMs per GPU

NICE DCV with External Rendering Server (KVM)

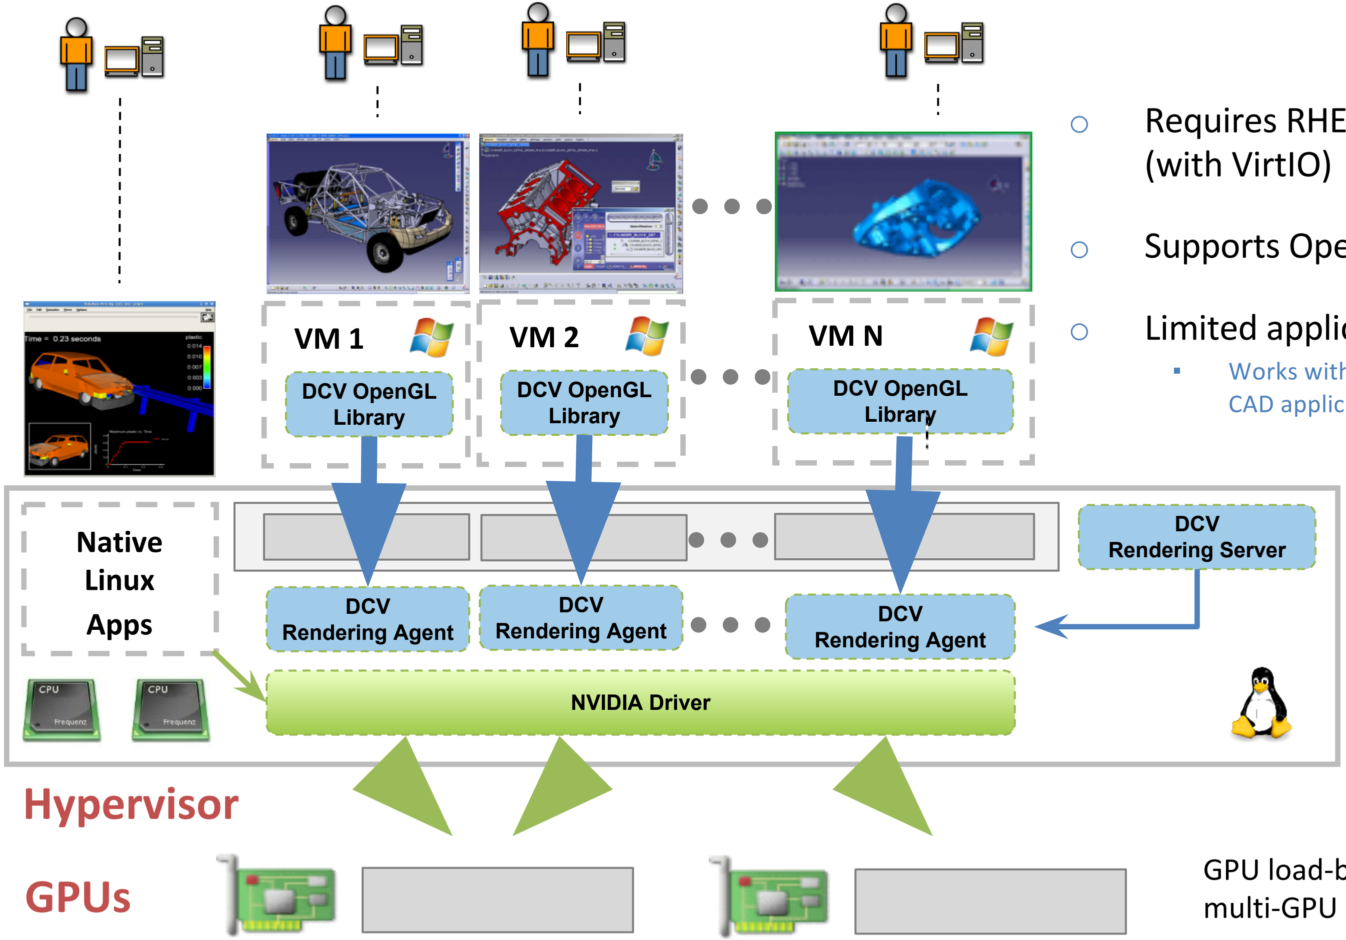

- Requires RHEL 6.x and above KVM (with VirtIO)
- Supports OpenGL 2.x, 3.x and 4.x
- Limited application compatibility
  - Works with most of the major CAE and CAD applications

Number of VMs limited by CPU, GPU and system memory

Native Linux applications can run on the host sharing the GPU with Windows VMs

GPU load-balancing on multi-GPU hosts

External Rendering Server: GPU appliance mode

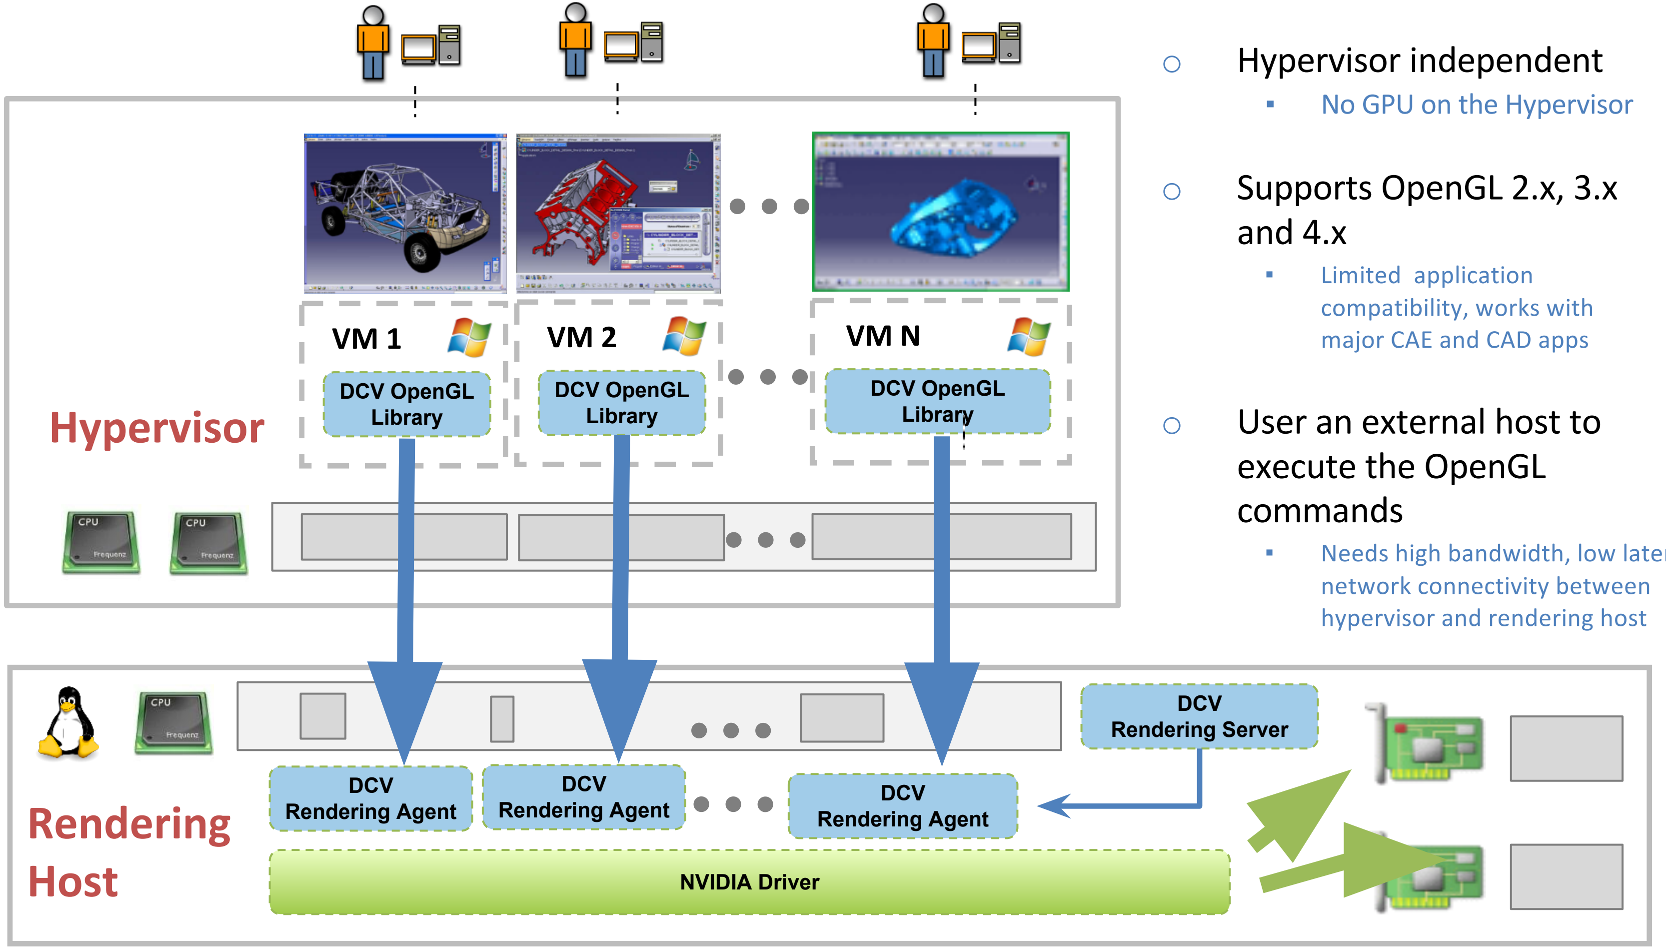

- Hypervisor independent
  - No GPU on the Hypervisor
- Supports OpenGL 2.x, 3.x and 4.x
  - Limited application compatibility, works with major CAE and CAD apps
- User an external host to execute the OpenGL commands
  - Needs high bandwidth, low latency network connectivity between hypervisor and rendering host

High Level Architecture (External Rendering Server)

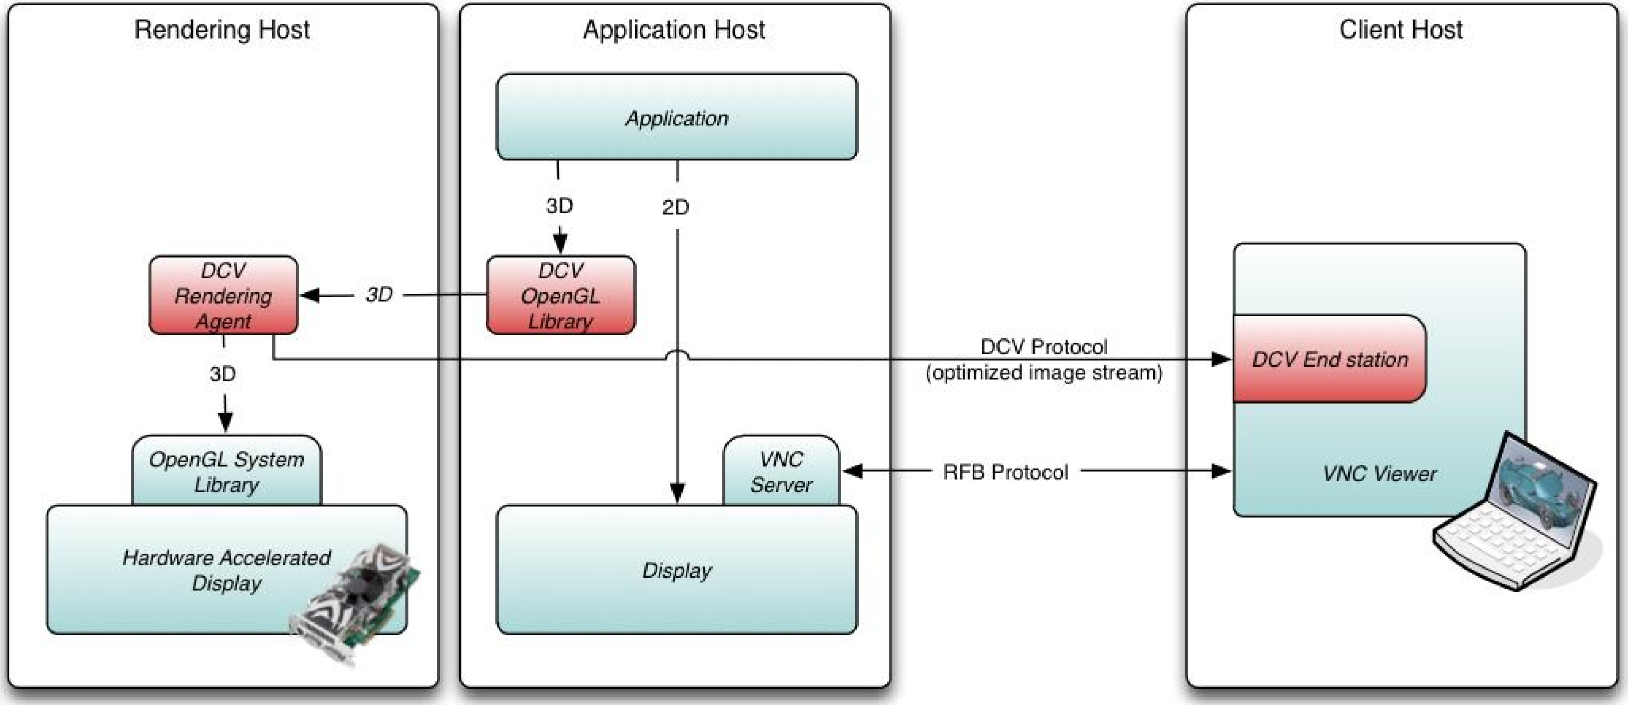

NICE DCV 2014.0 on VMs: Mode of Operation

Technology	Hypervisors	Application compatibility	OS Support	Pros	Limitations
GPU pass-through	All	Maximum	Linux and Windows	• Best performance	One VM per GPU
NICE External Rendering Server	All	Limited	Windows	• Best consolidation • GPU sharing • Dynamic GPU load balancing on multi-GPU hosts • Support for GPU appliance mode	No support for DirectX applications
NVIDIA vGPU	XenServer 6.2 SP1+ VMware vSphere ESXi 6.0+	Excellent	Windows	• Good performance • GPU sharing	Requires NVIDIA GRID cards and a specific hypervisor

Thank You

Thank you for listening… any question?

NICE Global Headquarter

Via Milliavacca, 9
14100 Asti (AT)
ITALY

Phone: +39 0141 90.15.16

NICE R&D Office in Sardinia

Edificio 1 Loc. Piscinamanna
09010 Pula (CA)
ITALY

Phone: +39 070 9243.2612

NICE USA Headquarter

2500 Citywest Blvd - S. 300
Houston, TX - 77042
USA

Phone: +1 (832) 699-0110

www.nice-software.com