



# Technical Visualization Workload Optimization

*Combined technologies from  
Adaptive Computing and NICE  
Software are making it happen*

## Running 3D Simulation the Right Way

Visualize this: You're a car designer and you're scanning a 3D simulation of your latest concept vehicle. You pan, zoom and otherwise manipulate the simulation that uses streamlines around the car to illustrate airflow. The simulation is massively pixelated and highly compute-intensive, and yet you're viewing and manipulating it on a low-end, three-year-old laptop. What's more, several of your colleagues are on a conference call with you, and they're manipulating those same pixels in the same session from multiple locations using typical office PCs and even mobile devices.

How is this possible? The answer is technical visualization workload optimization, which takes place in a highly efficient, secure, centralized location: the corporate data center. In fact, thanks to advanced technologies and the collaborative efforts of Adaptive Computing and Italy's NICE Software (see Figure 1), it is now possible to build a full-blown, physical-server-based Private Technical Compute Cloud. What Software as a Service did for 2D applications—keeping applications, compute resources and data together in the data center where they can be accessed by basic PCs, thin clients and mobile devices—can now be accomplished with complex 3D technical simulations. What's more, users can collaborate easily—anytime, anywhere—on the same session. And compute resources, including GPUs, can be used more efficiently than ever before.

Of course this scenario is in stark contrast to the status quo, in which technical visualization takes place largely at the desktop where expensive, high-powered Linux and Windows workstations sit side by side under users' desks, often idle or underutilized. The key technical visualization components—such as Fluent, Abaqus, CATIA, and other CAE/CAD applications—reside there too, and that's where they stay—unless collaboration is required, which brings up a whole host of other issues.

## Challenges Related to the Current 3D Model

Traditional pre- and post-processing of 3D modeling data in distributed computing environments has worked reasonably well for two decades or more. Generally, a user submits a simulation job to a cluster, where it is processed and the results are transferred back to the end user's workstation. Of course jobs might stack up and delays can be fairly common, but sooner or later output is sent down the wire and the data is visualized and interpreted.

However, when users need to collaborate, data needs to make the rounds, whether to a conference room across the hall or to users at remote facilities across town or around the world. So, networks are involved, and bandwidth and latency can become issues. This is true now more than ever, as average 3D simulation workloads continue to grow in size, with models in the 10 to 50 GB range being commonplace.

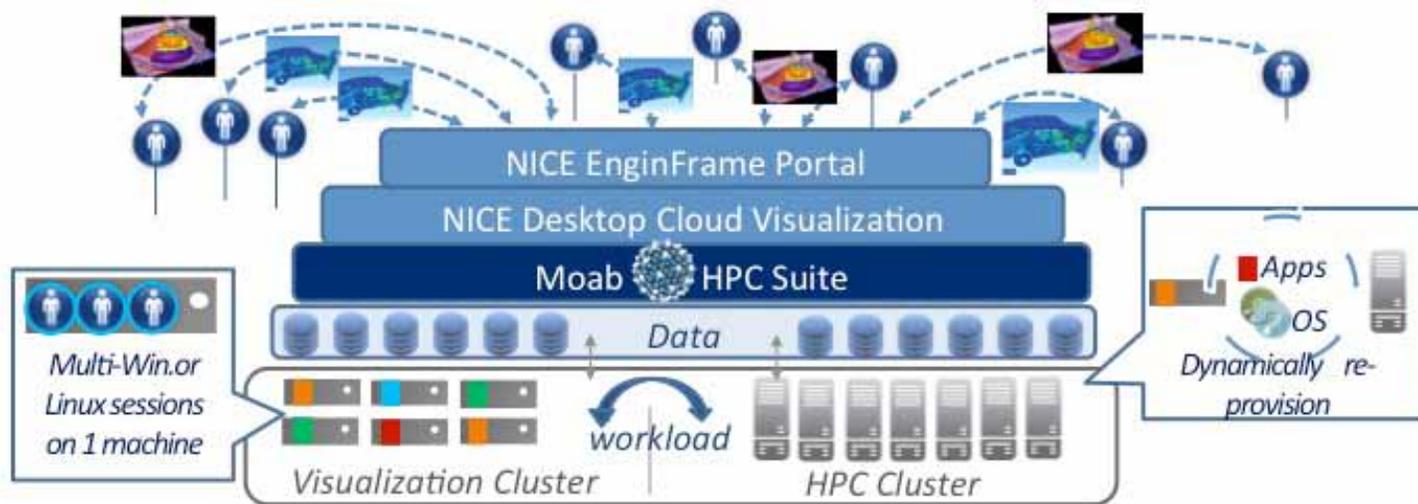


Figure 1. Moab HPC Suite works in concert with the NICE Desktop Cloud Visualization and NICE EnginFrame products to optimize technical visualization workloads.

Cost is also a key concern. Desk-side workstations usually feature high-end CPUs, top-of-the-line GPUs, and lots of memory. They have a short lifecycle—usually up to three years, but rarely more. And they are single-purpose machines. Once a simulation is up and running, the system’s usefulness is tapped out. Adding insult to injury, workstations must be sized for the largest potential simulation job. So, even if small or mid-sized technical visualization sessions are the norm, the workstation must be built to handle the worst-case scenario: that is, the most compute-intensive session imaginable. And that determination is made on the day the workstation is ordered. Two or three years down the road, IT is often dealing with unhappy users whose machines just aren’t powerful enough to handle increasingly CPU-, GPU-, I/O-, and memory-intensive applications and simulations.

While the up-front costs of desk-side hardware are substantial, ongoing expenses related to management of widely dispersed hardware and software are even bigger concerns in IT departments. Individual workstations have to be supported, updated, and eventually replaced, requiring expensive, time-consuming desk-side visits. They pose security risks, too. Not only do you have end users capable of going rogue and adding harmful applications to their systems, there’s the potential for leakage of proprietary data from multiple locations. And when confidential 3D simulation data is distributed in order to accommodate collaboration, the company’s risks can rise exponentially. Moreover, given today’s widely dispersed and mobile workforce, the two-workstations-under-the-desk paradigm doesn’t make nearly as much sense as it used to.

## **Moving 3D Applications Closer to the Data**

Rather than moving technical visualization data to authorized people wherever they might be and taking your chances, the logical solution is to move 3D applications closer to the data. Instead of having expensive hot, loud, underutilized workstations taking up space under desks, it makes sense to have an installed base of typical business PCs and laptops, with 3D CAD/CAE applications and data housed in a centralized, secure location: your data center. The applications can then be sent over the network to those basic desktop systems—in effect, moving pixels instead of data. Users get the full power of a high-end, GPU-enabled machine—as if it’s under their desks. But, in fact, the GPU and other high-performance compute

resources are housed on the server side and able to be shared by multiple sessions.

Being able to run multiple technical visualization sessions on a single server is a key benefit, but it’s only one of many. Compute capabilities can be “right-sized” on the fly, and utilization can be maximized using Moab resource management functionality. There is less network congestion and fewer bottlenecks, as moving visualized pixels to users results in a significantly smaller network burden compared to sending full data sets down the wire. Centralized and shared servers are more efficiently utilized than workstations. Support, updating, and replacement of hardware and software are more efficient, less costly processes, and they do not affect users nearly as much as when they are completed at individuals’ desks. Data stays in the data center and access control can be much tighter. Users are liberated and able to walk away from their desks and securely initiate a session, or log on to an existing one, in conference rooms, at home after work, or wherever they are authorized to do so from just about any device. In fact, workforce collaboration and productivity can improve dramatically. And remote, full 3D technical visualization and rendering can become standard capabilities of your organization.

## **Key Components of the Technical Visualization Private Cloud**

Remote, full 3D technical visualization is accomplished using key components of an integrated software solution developed through a partnership between Adaptive Computing, a cloud management software solutions provider and NICE, a visualization software and services company. Those components are:

**EnginFrame Views**, the customizable Web application portal that provides complete Linux desktop or single application views and enables access to the EnginFrame session-management environment for HPC and Cloud environments. Through the EnginFrame Views portal, the user can access a server-based 3D visualization application running on Linux or Windows, as well as the resources required to run the application. (A compatibility mode for iPad and Android-based tablets will be released later in 2012). In addition, data center administrators can provide encrypted and controllable access to remote users while monitoring technical visualization workflows in the data center—no HTML or Java knowledge required. IT infrastructure is protected, as is intellectual property, and secure, hassle-free collaboration is assured.

**Moab® HPC Suite**, the intelligence engine that maximizes resource utilization in technical visualization environments as well as conventional HPC clusters. Moab automates placement, scheduling, provisioning, SLA balancing and uptime for workloads based on multi-dimensional policies that mimic real-world decision making. It automatically balances the mission-critical workload priorities of enterprise HPC and technical visualization environments. Specifically with regard to 3D modeling and visualization, when the data center receives a session request, Moab assesses the size and application requirements of the request and the availability of data center resources—applications, CPUs, memory, GPUs, and nodes—and uses policies to allocate and schedule accordingly. It's sort of like playing Tetris with sessions and workloads—packing them into a grid as efficiently as possible. With a single click from a user requesting a VNC session, Moab snaps into action, determining the session requirements and scheduling the optimal session..

For organizations seeking to virtualize and dynamically provision both Linux and Windows resources, **Moab Cloud Suite** provides the framework for creating agile, automated and adaptive private or hybrid Clouds from diverse IT resources. It can manage setting up multiple Windows or Linux user sessions on a single machine to further maximize utilization. What's more, it can manage the dynamic re-provisioning of the OS and applications on a compute machine to better meet workload demand and maximize resource utilization and availability.

Moab can also optimize the allocation of storage and application licenses to allocate them from a shared pool across multiple users to reduce costs and improve utilization. Data storage can be consolidated into common, centralized storage nodes to reduce storage costs, and you can also reduce software costs with a shared common pool of applications that are dynamically allocated and re-allocated across multiple users as needed instead of underutilized individual licenses.

**TORQUE**, the open source resource manager has been aggressively developed and enhanced by Adaptive Computing and the open source community and is included in the Moab HPC Suite. TORQUE is designed to improve overall utilization, scheduling, and administration on clusters. In a Private Technical Compute Cloud, its primary function is to start sessions on the nodes Moab has chosen and keep Moab apprised of the health of the node and session. Together, Moab and TORQUE provide a highly intelligent resource orchestration platform.

**Desktop Cloud Visualization (DCV)** from NICE is the first fully accelerated remote 3D visualization product to enable GPU-sharing across multiple Linux and Windows desktops. One of DCV's key benefits is that it can aggressively compress simulation data in order to accommodate situations where bandwidth is in short supply. This is especially beneficial for users logging in from wireless or other relatively bandwidth-constrained network connections. DCV can toggle frames per second and modify the quality of the image depending on the quality of the connection and the user's level of interaction. When users move the mouse they get a high frame rate; by stopping the mouse they get a high-resolution image. Collaborating is as easy as clicking on a check box to allow access to a colleague, after which both users can manipulate the image as they see fit. DCV also offers a client-side application that lets users manually toggle image quality and shows bandwidth usage and frames per second. If network connectivity declines, fluid motion can still be maintained (at fewer FPS) by dialing back the desired quality. Installation of DCV is easy—point and click using just one CD.

DCV works with Red Hat's KVM-based hypervisor or other hypervisors on Linux servers, which enables IT administrators to run multiple Windows virtual machines on the same node (see figure 2). When a DCV driver is installed in a Windows VM, all OpenGL calls are captured and sent to the Linux OS for rendering, thus providing full GPU-sharing and acceleration to the Windows-based application, as well as GPU-sharing across multiple Windows OSes. So, for example, a Linux server containing three GPUs can typically support 10-12 Linux and Windows users, which certainly compares favorably to the one Windows workstation/one GPU model. Extending existing support for Linux and Windows running on physical nodes, DCV is now bringing full GPU acceleration to virtual machines, thereby enabling companies and Cloud providers to deliver virtual workstation services to their technical visualization users.

**VNC Visualization Edition**, a remote viewer that is used to access the VNC server. It works with both Linux and Windows and features view-only or full control of scalable remote desktop views. VNC Visualization Edition uses minimal system resources and is easy to install and use. When a session request is scheduled by Moab, a VNC viewer session is initiated on the user's endstation.

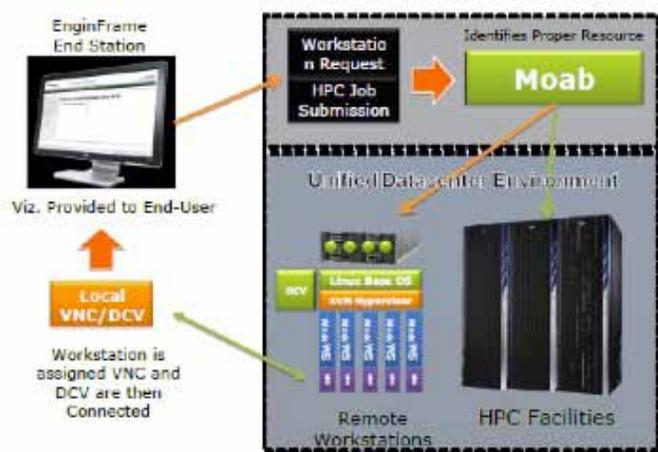


Figure 2. Moab and DCV, in conjunction with KVM or another hypervisor, enable Windows-based virtual machines to run on a Linux server and enable multiple Windows sessions on the same node.

## Visualization Workload Optimization Solution Highlights

### EnginFrame Views

- Web-based, user-friendly access to global visualization resources
- Moab integration-enhancing allocation
- Session sharing through portal
- Application session templates

### Moab HPC Suite

- Automation of session allocation, including multiple sessions on a single machine
- Enables HPC job submission and fair-share/SLA balancing
- Reserves resources for future scheduled sessions
- Applies policies to maximize resource utilization
- Controls number of users based on GPU capacities
- Adds priority rating to privileged user requests
- Dynamic re-provisioning of OS and applications on compute resources
- TORQUE manages resources to improve overall utilization, scheduling and administration on clusters
- TORQUE starts sessions and keeps Moab apprised of the health of the node and the job

### Desktop Cloud Visualization (DCV)

- Remote 3D visualization technology
- OpenGL virtualization layer

- Capable of multiple sessions on a single GPU
- Supports Linux and Windows
- Compensation for Low-bandwidth
- Collaboration-enabled

### VNC Visualization Edition

- Remote viewer for accessing VNC server
- Works with Linux and Windows
- Features view-only or full control of scalable remote desktop views
- Provides access to server sessions from clients of all kinds

## Visualization, Virtualization, and Consolidation

One of the biggest advantages of a Private Technical Compute Cloud is that multiple visualization sessions can run on a single Linux machine on which several GPUs reside. Once the users' applications are requested through EnginFrame Views, Moab determines the compute resources that are required, and TORQUE schedules the session. Many GPUs can be deployed on a single Linux cluster, and each of those GPUs can support multiple simultaneous sessions.

With Windows servers, on the other hand, session sharing isn't natively possible. As a result, Windows servers require one GPU per session, so a Technical Visualization Cloud with Windows servers has inherent limitations. This is less of a drawback than you might think, as Adaptive Computing has extended the capabilities of TORQUE to support Windows XP and Windows 7 for this solution. Moreover, DCV works with KVM or other hypervisors on Linux servers, which enables IT administrators to run Windows virtual machines on the same node (see figure 1). When a DCV driver is installed in a Windows VM, all OpenGL calls are captured and sent to the Linux OS for rendering, thus providing full GPU-sharing and acceleration to the Windows application, as well as sharing across multiple Windows OSes.

Moab Cloud Suite takes this functionality a giant step further by assessing the sizes and capacities of the various hypervisors on Linux servers in the data center and sizing and dynamically provisioning Windows-based virtual machines for sessions according to the application requirements. Thus, a Private Technical Compute Cloud environment can possess complete multi-OS /multi-session functionality.

## HPC Convergence

Many companies have major investments in HPC in the data center, and, generally, those resources aren't used at anywhere near capacity. The same can be said for technical visualization workstations that are dispersed throughout many companies. Therefore, the logical next step for many companies is to bring together those resources in a massive Private Technical Compute Cloud that handles both HPC and visualization. Depending on the types of GPUs installed, it is possible to share these resources between computational and visualization workloads. For example, if visualization systems are only used during the working day then these expensive devices can be used for overnight computational requirements. Moab can automatically handle this transition. Moab policies can be put in place to automatically determine which sessions and jobs go where. Those policies can also be set up to establish reservations for higher-priority workloads at specified times of the day or night. Moab policies can also give higher priority to certain privileged users—supervisors, for example—for sessions when resources are in short supply.

In addition, Moab has the ability to enable systems to be dual-booted. That is, a Windows OS can be swapped out for Linux or vice-versa. This can prove to be a valuable option when a Windows or Linux session is requested but the resources are not readily available. It brings a new level of flexibility to the data center.

## A Centralized, One-Stop Computing Shop Is Now Possible

Linux and Windows. Technical visualization and HPC. Distributed computing and centralization. Accessibility and security. These are polarities that present real challenges for IT management—challenges that keep a

lot of managers up at night. However, the good news is that technical visualization workload optimization, and the tools from Adaptive Computing and NICE Software that make it possible, can mitigate these challenges. Linux and Windows don't have to be isolated on disparate machines under users' desks. Technical visualization and HPC capabilities can be fully integrated in the data center. Users who are used to having their own high-powered workstations and applications at their fingertips can have a similar—and often superior—experience when compute power and applications are transferred to a central location, plus they gain mobility. With proper authorization, access can be had anytime, anywhere, and it doesn't have to be at the expense of security. In fact, by centralizing data, applications, and other key compute resources, security can be greatly enhanced.

Equally important, by consolidating resources in the data center for both technical visualization and HPC, utilization of resources and ROI are optimized, and IT support costs are minimized.

Technical visualization workload optimization provides the opportunity to deliver a higher level of flexibility, efficiency, security, and even frugality to your technical visualization support efforts. To learn more, contact Adaptive Computing today.

### Let's Talk . . . Set Up a Demonstration . . . and Test in Your Environment

- An Adaptive Computing solutions advisor can guide you to the products and services that will best meet your needs, and will work with you to set up a live, online interactive demonstration designed specifically for your organization.
- Contact a solutions advisor by phone or email, or visit our Web site today:

North America, Latin America . . . . . +1 (801) 717.3700  
Europe, Middle East, Africa . . . . . +44 (0) 1483 243578  
Asia, Pacific, Japan, India . . . . . +65 6597-7053  
Email . . . . . solutions@adaptivecomputing.com  
Web site . . . . . www.adaptivecomputing.com