

Implementing High Performance Workstation Virtualization for Data Processing in a Multi-User Microscopy Facility.

Daniel. E. Huber¹, John. M. Sosa¹, Jim Riedel², Chas D. Ellerbrock² and D. M. McComb¹

¹ Center for Electron Microscopy and Analysis, Materials Science and Engineering, The Ohio State University, Columbus, OH, U.S.A.

² Engineering Technology Services, College of Engineering, The Ohio State University, Columbus, OH, U.S.A.

Modern microscopy and microanalysis techniques are generating ever-larger data sets and require expensive computing resources for accessing, storing, archiving and processing of the data sets. This presents problem in a multiuser microscopy facility as researchers expect access to open view and process data collected on facility equipment with out needing to purchase a powerful graphics processing workstations especially. Currently the Center for Electron Microscopy and Analysis (CEMAS) at The Ohio State University has attempted to solve this problem through a combination of workstation and data center virtualization. The following paper seeks to address the issues relevant to facility managers and directors in the microscopy community. Considerations will be given to, performance, accessibility, security and cost.

Building on the efforts to provide the highest quality physical environment for electron microscopes, Figure 1, CEMAS has sought to optimize of workflows for data management and access to processing workstations. The common tasks at CEMAS, similar to many others, that require access to, or benefit from, high performance computing are image processing, 3D visualization, EDS map quantification and STEM image simulation. Each task benefits from a combination of machine capabilities, this is understood. However, access to large amounts of RAM memory, GPU based acceleration, I/O throughput and CPU compute speed benefits each respectively. The first and second tasks will be addressed in greater detail as they benefit from multiple aspects of the virtualization environment.

Currently the industry leaders in commercially available virtualization platforms are VMware, Citrix and Microsoft's HyperV. Virtualization has been available if not common place for servers and "back office" IT functions. Products such as Nvidia Grid K1 and K2 hardware have accelerated the ability to virtualize a graphics processing workstation by removing the visual lag and add support for GPU drivers lacking in prior virtualization products. While client "viewers" are now available for many platforms ranging from phones and tablets to traditional PCs and dedicated terminal hardware. CEMAS has evaluated all three platforms and currently implements VMware and HyperV as part of facility core IT infrastructure. It should be noted that the various different platforms support OpenGL, CUDA and DirectX drivers to different degrees and care should be taken before choosing a specific platform.

In summary, the centralization of computational resources can provide a means for the implementation of physical security and centralized management of workstation resources. While the benefits of centralized storage can provide acceleration of data intensive workflows. A successful strategy thus must balance the proper investment in hardware with the added complexity and additional IT expertise needed in implementing the virtual environment.



Figure 1. FEI Titan 80-300 installed at CEMAS at The Ohio State University.