

November 12, 2015 | By Jan Zverina

## Open Science Grid Using SDSC's Comet Supercomputer Virtual Clusters

### Successful Integration Seen as a “Milestone” for the Greater Research Community

The [Open Science Grid](#), a multi-disciplinary research partnership specializing in high-throughput computational services funded by the U.S. Department of Energy and the National Science Foundation, has added high-performance virtualized clusters to its global infrastructure by taking advantage of a new and unique capability of *Comet*, the National Science Foundation's newest supercomputer at the San Diego Supercomputer Center (SDSC).

The integration of *Comet* into the OSG provisioning system was led by a team including UC San Diego Professor Frank Würthwein, an expert in experimental particle physics and advanced computation. Würthwein joined SDSC, an Organized Research Unit of UC San Diego, in January 2015 to help implement a high-capacity data cyberinfrastructure across all UC campuses, as well as connect to key cyberinfrastructure organizations such as OSG. Würthwein was OSG's founding executive during 2005, and has again served as its executive director since February 2015.

*Comet*'s 'bare metal'-like approach means that a virtual cluster looks, feels, and performs almost exactly like the physical hardware. This enabled OSG to dynamically turn servers provisioned by *Comet* into [an HTCondor pool](#) and add new capability with very little additional overhead and significantly reduced administrative burden.

“Everybody wins in this collaboration, as OSG members are already conducting scientific research on this expanded infrastructure,” said Würthwein. “OSG's user community across physics, chemistry, biology, mathematics, and the social sciences gain transparent access to new capabilities, and neither SDSC nor OSG system engineers need to maintain a large new set of services that they wouldn't be supporting anyway.”

“Together, we’re creating a seamless interface between the nation’s two leading open scientific computing infrastructures – OSG and XSEDE,” said SDSC Director Michael Norman. “This latest effort is a major milestone for both SDSC and the OSG, as well as the entire research community. Frank’s additional role as a member of SDSC’s executive team enables SDSC and OSG to work together in pioneering advances in both high-performance and high-throughput computing.”

*Comet*, the result of an NSF grant worth almost \$24 million including hardware and operating funds, will be the first XSEDE production system to support high-performance virtualization at the multi-node cluster level. The cluster’s use of Single Root I/O Virtualization (SR-IOV) means researchers can use their own software environment, as they do with cloud computing, but achieve the high performance they expect from a supercomputer.

“We are pioneering the area of virtualized clusters, specifically with SR-IOV,” said Philip Papadopoulos, SDSC’s chief technical officer. “This will allow virtual sub-clusters to run applications over InfiniBand at near-native speeds – and that marks a huge step forward in HPC virtualization. In fact, a key part of this is virtualization for customized software stacks, which will lower the entry barrier for a wide range of researchers by letting them project an environment they already know onto *Comet*.”

Beginning with the next XSEDE allocation review in December, it will be possible to request allocations transparently across *Comet* and OSG.

“Scientists at campuses across the nation will be able to transparently compute from their desktops, labs, and campus infrastructures onto *Comet*, significantly expanding the reach of our new cluster toward what’s called the ‘long tail’ of science, or the idea that the large number of modest-sized computationally-based research projects represents, in aggregate, a tremendous amount of research that can yield scientific discovery,” said SDSC’s Norman. “We’re already seeing interest in *Comet*’s virtual clusters from other institutions, and expect that additional projects will enter production with them in the coming months.”

SDSC will present additional details about *Comet*’s high-performance virtualization at Supercomputing 2015 (SC15) in Austin, TX, November 16-19. Please visit SDSC at the SC15 exhibitor hall in booth #823.

---

#### MEDIA CONTACT

**Jan Zverina**, 858-534-5111, [jzverina@sdsc.edu](mailto:jzverina@sdsc.edu)

UC San Diego's Studio Ten 300 offers radio and television connections for media interviews with our faculty, which can be coordinated via [studio@ucsd.edu](mailto:studio@ucsd.edu). To connect with a UC San Diego faculty expert on relevant issues and trending news stories, visit <https://ucsdnews.ucsd.edu/media-resources/faculty-experts>.