

HUBzero + Cloud: Power tool for the masses

Dave Braun, Christopher Thompson, Alex Younts, Preston Smith and Carol Song

Rosen Center for Advanced Computing

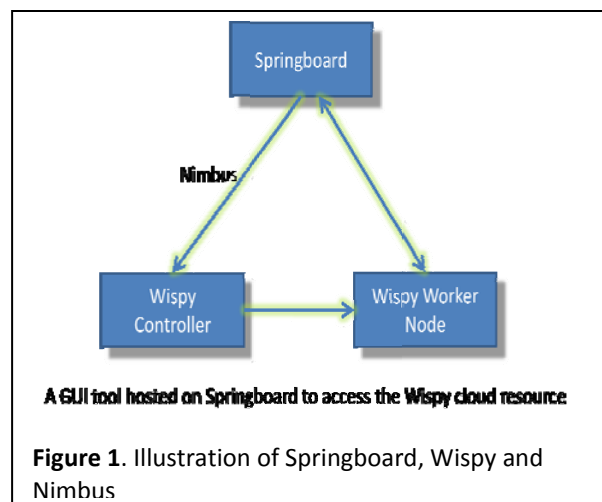
Purdue University

Today's powerful cyberinfrastructure, including the high performance and distributed computing systems and clouds, can be an unforgiving place for the non power users. Users are often expected to be familiar with Unix command lines, whether running computation jobs or moving data from one place to another. When they need immediate help, users typically rely on emailing technical support and system administrators, while in the commercial service provider world, online interactive help is fairly common.

How do we bridge these gaps and help scientific users to be more productive? We have developed *Springboard*, a web platform based on the successful HUBzero software, to provide easy access to HPC, cloud and other digital resources and online collaboration capability for instant help. Springboard has been integrated with the TeraGrid, including access to all the TeraGrid resources and relevant information (e.g., usage, allocation, etc). Leveraging the HUBzero middleware, especially the use of VNC (Virtual Network Computing), Springboard users are able to share their online sessions, including running of tools and applications, as well as login and command windows to the computing resources. This capability not only allows real-time collaboration among the collaborators, but also provides an online interactive platform that support staff can use to help their users.

This poster will feature a Springboard GUI application, **Cloud Dashboard**, that allows users to access cloud resources from Springboard in their web browsers through a GUI. The Cloud Dashboard connects from the Springboard web portal to Wispy, a cloud resource provided by Purdue University and a resource partner on the Science Cloud (<http://www.scienceclouds.org>). Wispy is a cluster running the KVM and Nimbus cloud software. It accepts user supplied virtual machine images and runs each VM with up to 4 CPUs and 16GB of memory. Users are able to package their applications and operating systems completely inside the images they submit and have full control over the execution environment. Wispy has been used to provide small, instant, on-demand clusters for various tasks and running complicated or prepackaged applications on additional hardware resources. There are currently 32 physical hosts (128 cores) on Wispy with a capacity limit of 128 VMs.

Cloud Dashboard supports both batch and interactive, or on-demand, types of applications. From its GUI, a user can load multiple VM images and run them on Wispy. The user may choose the length of runtime of the VM, whether to save the image upon termination, and with one click, can launch the VMs in the cloud. The VM status is reported back to the user in the GUI. While the VMs are running, the user may connect to the VMs via one button click and a separate login window is then displayed (Figure 2). The user may also choose to issue a command to be executed on a VM via the Dashboard (Figure 3). Integration with additional cloud resources is in the plan.



Tools such as the Cloud Dashboard have the potential to change the way HPC, cloud and other digital resources are typically used today, i.e., via command lines, which means a steep learning curve for a lot of people. These tools will enhance the usability of the HPC and cloud resources and allows the researchers to focus their attention and effort on scientific discovery rather than the intricate details of the computing and software environments. Together with the HUBzero's community participation and collaboration tools, Springboard hopes to bring the high end computing and data resources to the masses!

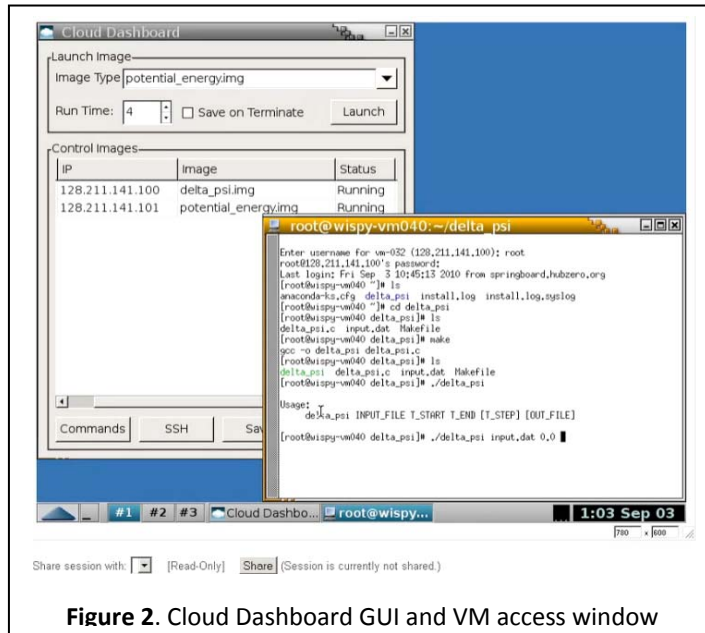


Figure 2. Cloud Dashboard GUI and VM access window

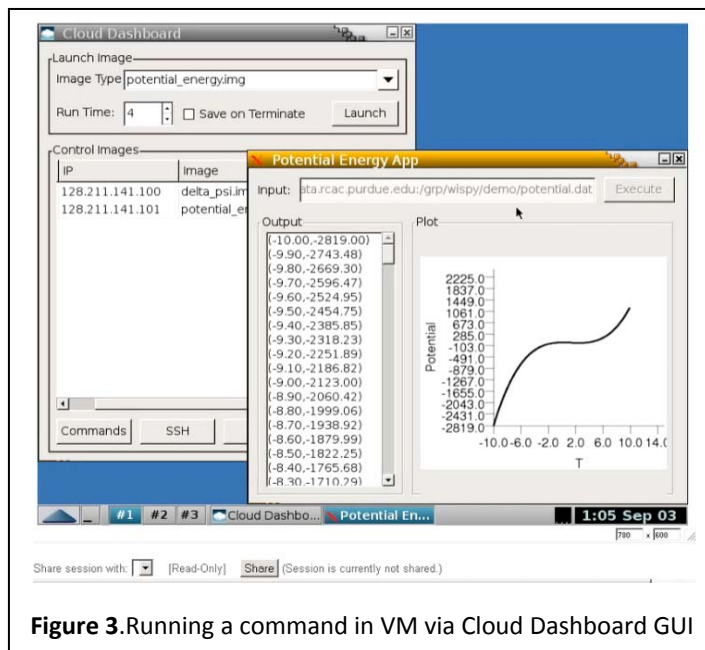


Figure 3. Running a command in VM via Cloud Dashboard GUI