VINO-Slurm: VIrtual NOdes in Slurm

I. Vardas[1#], **M. Ploumids**[1*] and **M. Marazakis**[1]

[1] ICS-FORTH

Presenting author: Ioannis Vardas, email: vardas@ics.forth.gr
* Corresponding author: Manolis Ploumidis, email: ploumid@ics.forth.gr

ABSTRACT

HPC systems are facing an increasing demand for resources from a wide range of workloads. Moreover, there is an increase in the type of workloads and thus their software-related requirements in HPC systems while isolation among different users is essential. Virtual Machine(VM) is a mature technology that allows both isolation among different users and privileged access to a fully customizable software stack with very low performance overhead. Additionally, HPC systems have multiples of concurrent users as well as a broad variety of resources. To allocate the resources and also manage the jobs issued by the users a special middleware is often employed, called Resource and Job Managing System (RJMS). One such tool, that is widely used in HPC is Slurm.

In this work we extend Slurm by adding support for running jobs in virtual nodes or else, virtual machines. Moreover, with VINO-Slurm users are allowed to choose the VM configuration (number of cores and memory) and also the VM image. The VM image allows the user to have privileged access to runtimes and libraries that are required by the application. The VM images are managed by VINO-Slurm which results in a simpler user interface. Finally, VINO-Slurm enables the users to run workloads either in virtual or physical hardware introducing the minimum possible differentiation in the way users interact with Slurm.