

Penguin Computing Offers HPC Compute Clouds Built for Academia and Research

2012-07-19

Penguin Computing today announced partnerships with multiple universities to enable easy, quick and unbureaucratic on-demand access to scalable HPC compute resources for academic researchers. "Penguin Computing has traditionally been very successful with HPC deployments in academic environments with widely varying workloads, many departments competing for resources and very limited budgets for capital expenses, a cloud based model for compute resources makes perfect sense," says Tom Coull, Senior Vice President and General Manager of Software and Services at Penguin Computing. "The new partnerships help academic institutions with a flexible cloud based resource allocation for their researchers. At the same time, they present an opportunity for IT departments to create an ongoing revenue stream by offering researchers from other schools access to their cloud." Penguin has implemented three versions of academic HPC clouds: Hybrid Clouds - Which are a local 'on-site' cluster configured to support the use of Penguin-on-Demand (POD) cloud resources as needed on a pay-as-you go basis. Local compute resources can be provisioned for average demand and utilization peaks can be offloaded transparently. This model lowers the initial capital expense and for temporary workload peaks excess cycles are provided cost effectively by Penguin's public HPC cloud. Examples of hybrid cloud deployments include the University of Delaware and Memphis University. Channel Partnership - Between Universities and Penguin Computing, allow educational institutions to become distributors for POD compute cycles. University departments with limited access to compute resources for research can use Penguin's virtual supercomputer on-demand and pay-as-they-go . Allowing them to use their IT budget for operational expenses. When departments I use is the university's HPC cloud revenue can supplement funding for IT staff or projects, increasing the department's capabilities. This model has been successfully implemented at the California Institute for Technology in conjunction with Penguin's PODshell, a web-service based solution that supports the submission and monitoring of HPC cloud compute jobs from any Linux system with internet connectivity. Combination Hybrid / Channel - The Benefits of the first two models have been successfully implemented at Indiana University (IU) as a public-private partnership. Penguin leverages the University's HPC facilities and human resources while IU benefits from fast access to local

compute resources and Penguin's HPC experience. IU can use POD resources and provide compute capacity to other academic institutions. The agreement between IU and Penguin also has the support of a group of founding user-partners including the University of Virginia, the University of California, Berkeley and the University of Michigan who along with IU will be users of the new service. The POD collocation offers access through the high-speed national research network internet2 and is integrated with the XSEDE infrastructure that enables scientists to transparently share computing resources. "This is a great example of a community cloud service," said Brad Wheeler, vice president for information technology and CIO at Indiana University. "By working together in a productive private-public partnership, we can achieve cost savings through larger scale while also ensuring security and managing the terms of service in the interests of researchers." For more information about Penguin Computing's HPC compute resources, please visit www.penguincomputing.com.