



CloudLaunch builds an elastic virtual HPC cluster in an Infrastructure-as-a-Service (IaaS) cloud. Using a cloud provisioning library and a set of SLURM plugins, CloudLaunch enables the SLURM scheduler to dynamically launch required compute nodes/cloud instances *on behalf of the job's submitter* using the submitter's cloud credentials, and tears down the instances when the job is completed. With CloudLaunch, HPC users can compute in an IaaS cloud using the familiar SLURM commands, without knowing the particulars of the cloud stack.

### **How does it work?**

The user submits a job, along with cloud credentials, to SLURM. SLURM is a popular scheduler and resource manager for HPC clusters, so it presents a familiar interface for running HPC jobs. CloudLaunch launches the required cloud instances using the user's credentials to run the job. When the job is completed, CloudLaunch terminates the instances. The user is billed according to the terms of the cloud provider for the usage of the nodes involved in the job. CloudLaunch supports multiple job queues, each with its own instance type and image/software so the administrator can configure the virtual cluster to run a variety of scientific applications.

### **What are the benefits of using CloudLaunch?**

- To individual users, CloudLaunch enables computing inside IaaS clouds using the familiar SLURM scheduler. The size of the virtual cluster is elastic depending on the demands of the submitted jobs, minimizing the wait time in queues during peak demand and the cost of maintaining idling infrastructure during low demand.
- To HPC administrators, CloudLaunch eases the management of compute nodes. Administrators can deploy new compute nodes and software in a new queue on demand by specifying a new node image and instance type.
- To cloud providers, CloudLaunch introduces cloud computing to a new class of users who may not have the expertise or resources to modify existing HPC software and scientific computing workflows to run in an IaaS environment.
- Because the instances launched by CloudLaunch are owned by the users, no special payment or billing mechanism is required. A user is billed simply according to the standard terms of the cloud provider whenever jobs are running.
- When used in a private IaaS cloud, CloudLaunch encourages sharing of resources within an organization by accommodating both HPC and non-HPC apps on the same set of hardware.
- CloudLaunch supports public clouds for maximum scalability.

### **What cloud providers does CloudLaunch support?**

CloudLaunch currently support AWS or Eucalyptus-based clouds such as Cornell's *Red Cloud*.

### **Where do I get more information?**

For more information, please contact the Cornell University Center for Advanced Computing at [help@cac.cornell.edu](mailto:help@cac.cornell.edu). The source code and documentation for CloudLaunch is not yet publicly available, but we do plan to publish it on GitHub soon.

