Contribution ID: **127** Type: **Sectional reports**

Evaluating Different Options for Scientific Computing in Public Clouds

Friday 9 July 2021 12:00 (15 minutes)

Cloud computing has emerged as a new paradigm for on-demand access to a wast pool of computing resources that provides a promising alternative to traditional on-premises resources. There are several advantages of using clouds for scientific computing. Clouds can significantly lower time-to-solution via quick resource provision, skipping the lengthy process of building a new cluster on-premises or avoiding long queue wait times on shared computing facilities. By providing a wide range of possible virtual machine configurations, clouds allow to easily adapt to changing workloads. Clouds can also reduce the total cost of ownership by allowing dynamic auto-scaling of computing resources depending on the current workload, or by leveraging spot instances that represent excess cloud capacity. A new serverless computing model has become popular recently, which enables users to seamlessly execute the so called cloud functions without having to manually manage and scale virtual machine instances.

Nowadays public clouds provide many options for running computing tasks ranging from manually managed on-demand virtual machines and HPC clusters to preemptible spot instances and cloud functions. This brings up several questions: which options are suitable for which kind of applications and use cases, what are their advantages and drawbacks, and how these options compare to traditional computing resources such as on-premises clusters. To answer these questions, we have implemented support for using the mentioned options as computing resources for running applications on Everest, a web-based distributed computing platform. This platform provides users with tools to publish and share computing applications as web services, and manages the execution of these applications on user-provided computing resources. Since Everest already supports the use of on-premises servers and clusters as such resources, this allowed us to evaluate and compare the new cloud-based resources against the traditional ones for execution of typical scientific computing applications such as bag-of-tasks and workflows. This approach also enables simple migration of existing applications to these new resources.

In this report we will describe the implementation of new cloud-based resources for Everest and will present the results of their experimental evaluation and comparison.

Summary

Primary authors: Dr SUKHOROSLOV, Oleg (IITP RAS, NRU HSE); Mr KRASNOV, Svyatoslav (NRU HSE); Mrs TSAREVA, Yulia (NRU HSE)

Presenter: Dr SUKHOROSLOV, Oleg (IITP RAS, NRU HSE)

Session Classification: Virtualization

Track Classification: 7. Virtualization