

Public cloud for high throughput computing

Friday 29 March 2019 09:50 (25 minutes)

The vast breadth and configuration possibilities of the public cloud offer intriguing opportunities for loosely coupled computing tasks. One such class of tasks is simply statistical in nature requiring many independent trials over the targeted phase space in order to converge on robust, fault tolerant and optimized designs. Our single threaded target application (50-200 MB) solves a stochastic non-linear integro-differential equation relevant for read/write simulations of heat assisted magnetic recording (HAMR) for high areal density hard disk drives (HDD). Here, the phase space is multi-dimensional in physical parameters and potential recording schemes. Furthermore, for any one such point in phase space, 100s of simulations must be repeated due to the stochastic nature of the physical simulation.

In this talk, we show that a simple abstraction layer between the target application and cloud vendor provided batch systems can be easily constructed thus avoiding changes to the underlying simulation and workflow. With some planning, this abstraction layer is portable between three available cloud providers: Amazon Web Services, Microsoft Azure and Google Cloud. This abstraction layer is required to be light weight and not introduce significant overhead and was implemented as simple Bash scripts. To reduce cost, it was critical to test the application under multiple configurations (e.g. instance types and compiling options), avoid local block storage and minimize network traffic. Fleets of 100,000 concurrent simulations are easily achieved with over 99.99% of the cost just for compute (versus storage or network). By implementing a third party grid engine, 1,000,000 concurrent simulations were achieved with no modifications to the abstraction layer.

Best practices and design principles for HTC in public cloud will be discussed with emphasis on robustness, cost and horizontal scale and unique challenges encountered in this migration.

Primary author: Dr PARKER, Gregory (Entonos)

Presenter: Dr PARKER, Gregory (Entonos)

Session Classification: Grid, Cloud and Virtualization

Track Classification: Grid, Cloud & Virtualisation