

CS554 Project Ideas

CloudKon:CKHPC – Extending CloudKon to support HPC Jobs Scheduling

Overview

Predictions are that by the end of this decade, we will have exascale system with millions of nodes and billions of threads of execution. Unfortunately, today's job schedulers have centralized Master/Slaves architecture (e.g. Slurm, Condor, PBS, SGE), where a centralized server is in charge of the resource provisioning and job execution. This architecture has worked well in modest scales and coarse granular workloads, but it has poor scalability at the extreme scales of petascale systems with fine-granular workloads. The goal of this project is to leverage the newly developed CloudKon system which supports distributed task management for Many-Task Computing workloads through various cloud services (SQS, DynamoDB, and EC2). This project is about adding support for HPC jobs in CloudKon, distinguishing MTC tasks and HPC jobs by the number of cores or nodes a task/job requires. In MTC, tasks are generally single core, or single node. In HPC, jobs are typically multi-core or multi-node, and hence extra coordination is needed to ensure that multiple cores or nodes are started and stopped at the same time. MPI applications should be supported. The solution should be compared to Slurm. This project requires 2~3 students.

Relevant Systems and Reading Material

CloudKon: http://www.cs.iit.edu/~iraicu/research/publications/2013_Qual-IIT_CloudKon.pdf

Amazon SQS:

- <http://aws.amazon.com/sqs/>
- http://sqs-public-images.s3.amazonaws.com/Building_Scalabale_EC2_applications_with_SQS2.pdf
- <http://awsdocs.s3.amazonaws.com/SQS/latest/sqs-gsg.pdf>

DynamoDB: <http://aws.amazon.com/dynamodb/>

SLURM: <https://e-reports-ext.llnl.gov/pdf/241220.pdf>

Methodology

This project will extend CloudKon to support multi-core and multi-node HPC jobs.

Preferred/Required Skills

Programming language choice: Java, C/C++, Python

Skills/knowledge: REST API, Amazon EC2, Amazon SQS, Amazon DynamoDB, Linux Bash Scripting, distributed queues, MPI.

Performance Metrics

Throughput, Latency, Efficiency, Utilization

Project Mentor

Iman Sadooghi: isadoogh@iit.edu, <http://datasys.cs.iit.edu/~isadooghi/>

Ke Wang, <http://datasys.cs.iit.edu/~kewang/>