

CS554 Project Ideas

Fabriq:Sched: Leveraging Fabriq as a building block for Distributed Scheduling with CloudKon+

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. CloudKon is a distributed task scheduler that leverages Amazon Simple Queuing Service (SQS) as a public cloud service to provide a scalable task scheduling system that supports Many Task Computing (MTC) workloads. SQS is a distributed queue service with the purpose of providing content delivery on extreme scales. Therefore there are some requirements for MTC type workloads that cannot be met by SQS. For example SQS guarantees at least once message delivery, but not exactly once delivery. CloudKon uses DynamoDB for that matter to get rid of the duplicate tasks.

One of the limitations of CloudKon is the fact that it uses SQS. Running on top of SQS, makes it impossible to run on other public or private cloud infrastructures. In order to overcome this limitation we have proposed Fabriq. Fabriq is a distributed message queue that runs on top of a Distributed Hash Table, namely ZHT. The design goal of Fabriq is to achieve scalability and near perfect load balancing, while handling the operations with minimal overhead and low latency. Moreover, Fabriq guarantees exactly once delivery. That means there will be no need for another component to eliminate duplicate tasks. The goal of this project is to build on top of Fabriq a distributed MTC scheduler (called CloudKon+), similar to that of CloudKon (with the exception that it would be platform independent). Fabriq is written in C++. In order to achieve the best performance, and for your system to run in most distributed systems environments, CloudKon+ needs to be written in C++.

Relevant Systems and Reading Material

- Fabriq paper: (Available on Blackboard)
- CloudKon: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6846476&tag=1, <https://sites.google.com/site/imansadooghi/cloudkon/how-to-use-cloudkon>
- Many Task Computing paper: I. Raicu, Y. Zhao, I. Foster, "Many-Task Computing for Grids and Supercomputers," 1st IEEE Workshop on Many-Task Computing on Grids and Supercomputers (MTAGS) 2008. http://datasys.cs.iit.edu/events/MTAGS08/MTAGS08_p25.pdf
- Sparrow: Scalable Scheduling for Sub-Second Parallel Jobs; <http://www.eecs.berkeley.edu/Pubs/TechRpts/2013/EECS-2013-29.pdf>
- SLURM paper: <https://e-reports-ext.llnl.gov/pdf/241220.pdf>

Methodology

First, you will have to analyze the architecture of the CloudKon and Fabriq. Then, you need to implement the CloudKon+ based on the logic use on CloudKon and Fabriq.

Preferred/Required Skills

Required: Linux, C++.

Evaluation and Metrics

Operation latency and its distribution, throughput, scalability, CDF graph. Metrics used in CloudKon paper. Testbed should be Amazon EC2, and scales should be up to 128 instances.

Project Mentor

Iman Sadooghi, isadoogh@iit.edu