Client

Daniel Avery, ICR davery@icr-team.com

Background

Our current application is being used to generate critical reports for our customers. It utilizes a microservice architecture with over 20 microservices running in a Kubernetes cluster. This application is being powered by a simple relational database powered by Amazon's RDS. The current database implementation needs to be upgraded for performance reasons for new use cases. We want to be able to improve basic query performance as well as set ourselves up for future work in AI/ML.

Project Goals and Requirements

You will explore a few different database options to provide critical performance information that we will use in our application. The goal is to determine which database is most performant with multi-dimensional, multi-relational data. Ideally, students will be able to characterize which database technologies work best for traditional queries and in terms of pulling data for use in AI/ML algorithms. This will include:

- Relational database
- Graph database
- NoSQL database

Students will be able to learn Docker, various databasing technologies, scripting, and if time allows some basic AI/ML algorithms. Additionally, students will be given the opportunity to present the results in a professional environment, learning the soft skills needed to be successful in the software industry.

Suggested Team Size:

4 students, US Citizens Preferred

Skills / Experience for CSM Students:

- Various database technologies (see above)
- Python or some other scripting language to insert data and AI/ML
- Java for database interface