

Cloud-based HPC Service for University Medical Research

Creating a high-performance computing environment to improve automation and scaling, optimizing costs and utilizing the power of CycleCloud orchestration.

OVERVIEW

The HPC Team at Oakwood had the opportunity to work with one of the nation's most prestigious higher-learning institutions to create a high-performance computing environment in the cloud.

The university's health and science division works in genomic testing and other research endeavors that require significant processing speed and accuracy. They rely on this data to further advance their knowledge and understanding in life sciences.

This new cloud based HPC environment will allow them to meet their needs of future research ambitions while moving away from aging on-premise infrastructure.

Technology needs...

- Uses groups of high-speed servers (clusters) with a centralized job scheduler.
- Works in parallel at extremely high speeds.
- Processes massive multi-dimensional datasets.

Typical use cases...

- Sequence DNA.
- Run artificial intelligence (AI) algorithms and simulations.
- Analyze terabytes of data streaming from IoT devices in real time.

CHALLENGE

The clients knowledge and understanding of HPC environments on-premise was solid but their understanding of how to build, scale, and optimize within the Azure cloud environment was limited.

SUMMARY

Business Challenge:

Limited skills in migrating on-premise HPC environment to cloud. Maintaining NIST-800 compliance policies.

Solution:

Creation of a new CycleCloud template for the client that included custom configurations for both Slurm and Lustre that accommodated, not only the requirements for throughput and storage needs but, being able to be NIST-800 compliant.

Benefits:

Ease of management and configuration when needed via CycleCloud's user interface while meeting the backend technical requirements of the overall solution.

There were many challenges related to the size and volume of data, including the number of small and big files any research lab could be working on at any given time.

The Team was small and limited with resources to accomplish this task of setting up a cloud-based solution as well as both resource and time constraints against them. Combine that with limited skillset around the Lustre parallel system and CycleCloud templates - there was a lot of work to be accomplished in order to be successful.

The security team at the client site was also rolling out NIST-800 compliance policies that conflicted with out of the box solutions that came with CycleCloud and Lustre marketplace image. This required many custom changes to accommodate the resources to become compliant with the policies being pushed out into the environment.

SOLUTION

Oakwood created a new CycleCloud template for the client that included custom configurations for both Slurm and Lustre that accommodated, not only the requirements for throughput and storage needs but, being able to be NIST-800 compliant.

This included a CycleCloud specific Lustre template, and a Slurm template that included the Lustre client installation. The result was ease of management and configuration when needed via CycleCloud's user interface while meeting the backend technical requirements of the overall solution.