

PROJECT DESCRIPTION

Our project explores the scalable containerization of VOLTAplayer with Kubernetes. It also explores the cost/benefit trade-off of leveraging elastic cloud resources versus High Performance Computing (HPC) to perform \ multidisciplinary design optimization (MDO) analysis.

What we did:

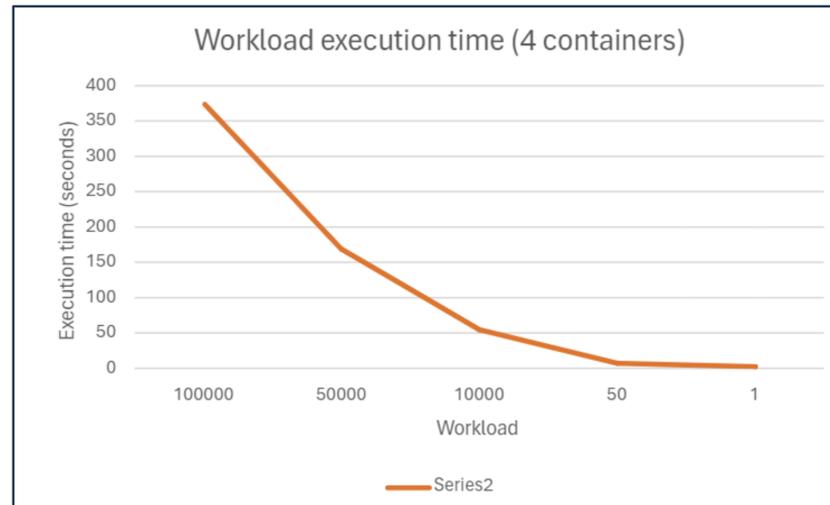
- Created Containerized instances of VOLTAplayer (through Kubernetes and Docker) to streamline testing process
- Gathered performance metrics of VOLTAplayer running on Geddes, one of Purdue University's in-house HPCs, as well as on elastic cloud computing resources
- Conducted a Cost-Benefit Analysis between the cloud computing resources as well as HPCs, taking into account maintenance costs, initial investments, as well as engineering turnaround time increases caused by slower MDO analyses.

Future Goals

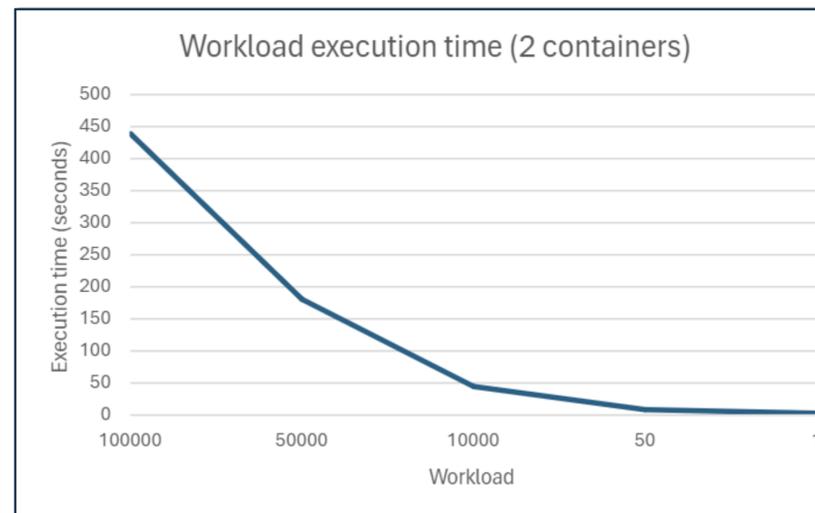
- Perform a comparative analysis between different cloud providers (AWS vs. Google Cloud Platform), and between different HPCs (Anvil vs Geddes). This is to identify which platforms cater best to the company's specific needs.
- Observe performance under high demand workloads and establish metrics of scalability for cloud providers vs. HPCs. When workloads increase, this will help determine what new resources need to be allocated.
- Explore methods to optimize resource utilization on HPCs to minimize workload execution times and efficiently schedule tasks. This will help minimize operation costs in the long run.

References

- The "Geddes" Compostable Platform <https://par.nsf.gov/servlets/purl/10216791>
- AWS Pricing Calculator <https://calculator.aws/#/>



Workload vs Execution Time on HPC using 4 pods



Workload vs Execution Time on HPC using 2 pods



HPC environment



Amazon Web Services Logo



Docker Logo



Kubernetes Logo

Research Methodology

To perform the cost benefit analysis, we established two distinct environments to run Volta.

- Container that could be deployed on a cloud service such as AWS (Amazon Web Services)
- Purdue Geddes HPC (High Performance Computing)

Allowed us to experience high computational power of Geddes and flexibility of containers.

Ran projects on each version of Volta and gathered performance data.

Researched operation costs using Purdue and AWS websites.

By comparing performance data and cost, we determined the optimal platform.

Conclusion

- AWS is more scalable and provides more flexibility but can't support as much computing power compared to an HPC.
- Hosting a HPC environment similar to Purdue Geddes would have more upfront/general costs but it would have improved speed and efficiency and would be more secure than using third-party cloud services.
- Limitations and biases include only testing with one file and only testing a couple workloads for each pod.

Acknowledgements

- Thank you to the Lockheed Martin and the ESTECO team for their continued support
- Thank you to Daryn Decker, Francis Raymundo, and Shaun Scott for their guidance throughout this year
- Thank you to the Data Mine team for their guidance and Khushi Saini our teaching assistant