



1001 Craig Road Suite 305
St. Louis, MO 63146
sales@oakwoodsys.com
www.oakwoodsys.com

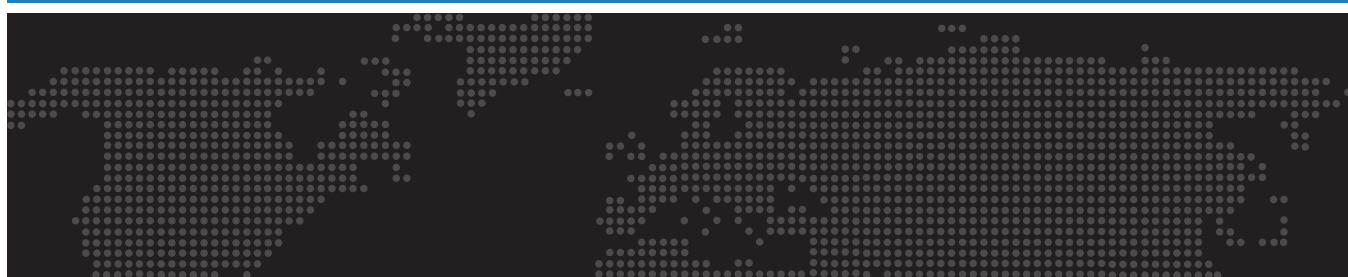



WHITEPAPER

HIGH-PERFORMANCE COMPUTING IN AZURE

Technologies Discussed:

Microsoft Azure Cloud
High-Performance Computing (HPC)
Azure CycleCloud
Lustre
Slurm





About Us

Headquartered in St. Louis, MO, Oakwood Systems Group is a Microsoft Gold Partner specializing in Cloud Migrations, Application Modernization, Data & Analytics, Security and Managed Services.



Since 1981 the Oakwood Team has been helping organizations of all sizes, across all industries, solve their most complex technology-related challenges. Our longevity has allowed us to evolve and sharpen our delivery approach(es) to help guarantee better business outcomes for our clients.




I am really happy about how the migration went, as is the entire team! We appreciate your approach and communication throughout the process. Thank you for the kudos to our team and we are sending the same your way! We also wish we worked with more vendors like you.

- Jennifer Walton | St. Louis Zoo


We see a future where our clients' success, advanced by our solutions, is so profound that they become our most powerful advocate. The Oakwood Team is driven to excel and is accountable in delivering quality results.





What We Do

The Oakwood Team is comprised of seasoned technology experts well-versed across the entire Microsoft stack. They understand how to best leverage technology to meet their client's goals.



Cloud & Infrastructure

Building optimal IT environments to support business transformation for organizations large and small.



Data & Analytics

Showing clients how to leverage critical business data to become an insights driven organization.



Application Development

Experts in building custom mobile, web, and business applications that deliver outstanding customer experiences.



Security

Helping protect client's business data and their users from cyber criminals looking to exploit security gaps in their environment.





Overview

The purpose of this paper is to provide an overview of a 2021 project involving Oakwood's Azure Team and their work with a world-leader in elite genetics and biotechnology serving farmers around the globe.



Recently, the Oakwood Team had the opportunity to work on an incredibly unique project involving a leader in elite genetics and biotechnology serving farmers around the globe. The data they provide to livestock owners is critical in helping them breed better quality animals while finding efficiencies in the production process.

This project involved working with the client in rethinking how to better deliver relevant and timely data to their customers through a high-performance computing environment.





Need 4 Speed

Improving the client's High-Performance Computing environment to deliver massive amounts of complex calculations quickly.


A Genetics Testing Computing Problem

Genetics is inherently complex and involves a massive amount of mathematical calculations. The client needed to ensure that these calculations are reliable and accurate to meet the demanding needs of their customers. Producing this data manually in a reasonable amount of time is simply not possible. As such, the client relies on computations that are performed by large computer clusters.

This genetic testing puts a lot of strain on the client's computer resources. The client had a high-performance computing (HPC) environment set up to handle their genetic testing requirements. Improving this environment would help farmers around the world breed better animals.

The client was looking for ways to improve their computing environment, as well as migrate some of their data to Azure. They needed to find an expert well-versed in cloud migrations that could handle complex data requirements.

Luckily, the genetic testing company was able to find the right partner in Oakwood Systems Group.





Partner

The client was searching for partners on Microsoft's Partner Network and Oakwood met the criteria of applicable gold competencies and decades of relevant experience.

Building A Partnership For A Brighter Future

Oakwood Systems is a Microsoft Partner holding 8 Gold Competencies and 6 Silver Competencies. With over 40 years in business combined with several Microsoft Partner of The Year awards under their belt – this client was confident they found the right partner with the depth and breadth of skills capable in delivering on their lofty initiatives.

The client reached out to Oakwood in 2021 and, after several conversations, the project kicked off in May 2021. Oakwood's Azure Team of seasoned consultants set out to redo the client's high-performance computing cluster, as well as work on data migration. There were several project requirements that needed to be met for this project to be successful.

The client required improved computing performance with higher availability, higher reliability, and more processing power. These requirements needed to be met while ensuring that the client's team would have access to the same, or similar, technologies that currently integrated within their existing systems.

Some of the client's data would have to be migrated to Azure while maintaining an appropriate security posture. Additionally, Oakwood had to ensure that the high compute clusters would operate effectively, and that the client had the necessary throughput for compute cycles in the Azure environment.

Unsurprisingly, there were issues that the Oakwood team would have to overcome.





Challenges

As with many projects - the need to adapt to challenges along the way requires a Team well-equipped to offer solutions to challenges as they arise.

01

Never Been Done Before

The documentation for this type of project was only theoretical. There was no real-world example of a successful implementation. The level of scalability had never been done before. Additionally, the client already had a custom-built environment to meet their needs. Transitioning that environment would be no easy task. This project would require a custom development solution as there was no cloud-native solution available from Microsoft.

The Oakwood team would have to develop new templates, new configurations, and create custom templates.

02

Proprietary Data

Due to the proprietary nature of the client's work, there were limitations to the data that they could share with the Oakwood Team. This lack of visibility created a number of roadblocks. The proprietary data created issues when developing the planned solution. Oakwood had to be sure that the client's users could only access data and systems that they needed for their day-to-day work. Giving users more access than they needed could create major security risks.

03

Moving Goal Posts

The project requirements and deliverables changed multiple times throughout the course of the project. For example, there were numerous changes to the disaster recovery process. Oakwood and the client went back and forth on how to automate disaster recovery to ensure consistent availability of the computational environment.





Implement

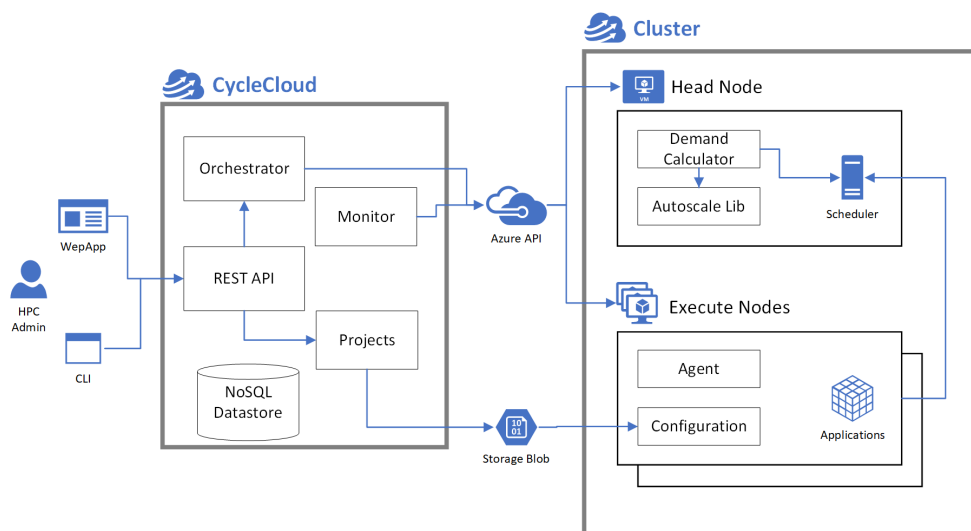
Turning to Azure CycleCloud for the operation of the client's High-Performance compute cluster environment.

Project Implementation

Even through the previously discussed difficulties, the Oakwood Team maintained a positive mindset and worked closely with the client to develop solutions and work-arounds. Oakwood knew that they could develop a comprehensive solution to help the client better meet their business needs. The Azure Cloud would be the backbone in helping the client to free up computer resources, increase redundancy, and provide reliable system backups. Oakwood needed to decide what tools would best assist the client in the operation of their high-performance compute cluster environment. The answer was found in Azure CycleCloud.

Azure CycleCloud provides an easy way to manage HPC workloads where users can simply create and manage various HPC clusters. Azure CycleCloud is ideal for deploying HPC schedulers, and automatically scaling the infrastructure for computational needs.

First, the team focused on researching CycleCloud and how it would work in the Azure environment. Oakwood had to ensure that the HPC environment would deploy with little or no involvement from the client's end. This required a fair amount of customization and tweaks to the user interface of the various tools they used.





Managing Data

Building a fast and on-demand file management system to provide necessary throughput for compute cycles.

Lustre Environment in CycleCloud

A proper file management system was needed to ensure that the client would have the necessary throughput for compute cycles. The Oakwood Team found that deploying a Lustre environment in CycleCloud would meet the client's needs. Lustre provides a scalable, parallel file system purpose-built for large HPC environments.

Oakwood was able to leverage Lustre to build a fast on-demand, high-performance file system. This system would give the client the necessary throughput for their compute cycles. Once the Team landed on Luster, they needed to determine how best to deploy the Lustre environment in CycleCloud. Additionally, storage space in Lustre was somewhat limited.

The backend of Lustre was managed using Azure storage and hierarchical storage management to lower overall usage. Oakwood used a Robinhood policy server to manage the disc and file usage. Robinhood is an open-source tool that assists in the management of large file systems.

Another piece that needed to be addressed was monitoring. The client needed to monitor the metrics of their HPC environment. Oakwood turned to Telegraf to address this need. This tool allows users to collect custom metrics in most environments and Oakwood was able to successfully configure this tool with Lustre.

Data Needs

The next step was to ensure proper data integration throughout the client's processes and systems. The Oakwood team settled on Slurm for the ingestion and transportation of data. Slurm is a highly configurable open-source workload manager that is already built into CycleCloud.

However, the standard template would not work within the client's unique environment. The Oakwood team heavily modified and configured the Slurm template and environment. Slurm was tied back into the on-premise servers. Additionally, dealing with updates in Slurm could be a potential roadblock. The Oakwood team had to make additional changes to ensure it would work for future versions.





Microsoft Team

In order to accomplish some of these modifications, Oakwood worked directly with the Microsoft HPC Team.

Oakwood Partners with Microsoft HPC Team

Oakwood was able to work directly with Microsoft specialists to overcome difficulties they ran into throughout this project. The Oakwood team worked with Microsoft to modify both Lustre and Slurm templates. Microsoft even dedicated their own resources in writing custom scripts with this project in mind.

The Microsoft specialists were more than happy to assist with this project. The specialists were excited to use their tools in a live environment. Oakwood team members would break the tools, such as Slurm, and then work with the Microsoft team to fix it. Microsoft was a key supporter in this project.





Outcome

Oakwood delivers performance and functionality improvements which now requires less day-to-day management.


Project Success

Oakwood's client has seen tremendous performance improvements and is still working with Oakwood on expanding the performance and functionality of this complex environment. The newly implemented system is more sustainable and workable than their on-premises systems. The client can automatically build and run entire HPC clusters to run genetic tests. Their team can achieve a two-gigabyte throughput using Lustre as backend storage.

The client has to perform less day-to-day management of their systems as compared to before.

Azure has freed up significant computer resources for the client. The client has more redundancy and backup than ever before. Overall, this customer is better able to meet the needs of their customers by using their new HPC environment.

The improvements that Oakwood was able to deliver is clearly apparent to the client. Oakwood plans to continue to work on business intelligence, as well as data and reporting functions. The client and Oakwood plan to use the lessons learned from this project on future high-performance computing clusters and data center deployments.



Microsoft Partner



Gold Cloud Platform
Gold Cloud Productivity
Gold Datacenter
Gold Windows and Devices
Gold Collaboration and Content

Microsoft Partner



Gold Application Development
Gold DevOps
Gold Data Analytics
Gold Data Platform
Gold Cloud Platform



1001 Craig Rd. Suite 305
St. Louis, MO 63146
marketing@oakwoodsys.com
www.oakwoodsys.com