

**SILICON
MECHANICS**



**VAN ANDEL
INSTITUTE**

HYBRID HPC-PRIVATE CLOUD DEPLOYMENT OVERVIEW

Executive Summary

Silicon Mechanics, a leading provider of cloud solutions and high-performance computing technologies to the world's most innovative organizations, was chosen by the Van Andel Institute (VAI), an independent biomedical research and science education organization, to design a hybrid HPC-private cloud solution that could provide massively parallel processing and analysis capabilities while allowing VAI users a choice between both traditional cluster and cloud resources.

This guide summarizes how Silicon Mechanics architected, built, and deployed a hybrid cloud HPC solution for VAI in under six months. The deployment process surpassed the expectations of the VAI leadership, and was widely promoted throughout the tight-knit life science research community as a noteworthy success case.

Hybrid HPC-Private Cloud Deployment Overview

Background

The Van Andel Institute (VAI), based in Grand Rapids, MI, is a nonprofit biomedical research and science education organization engaged in the investigation of the epigenetic, genetic, molecular and cellular origins of cancer, Parkinson's, and other diseases. The organization employs 330 scientists, educators and staff.

When VAI first reached out to Silicon Mechanics, they had virtually no ability to handle big data in-house. Large-scale scientific computing projects were either not being processed due to this lack of on-site infrastructure, or offloaded to third party organizations like the National Science Foundation (NSF). The lack of in-house compute capacity was hampering the speed and effectiveness of research, forcing scientists to use valuable research time to seek outside funding and resources, while also impacting the institute's recruitment of scientists who required terabyte- and petabyte-scale storage project capabilities.

VAI hired research computing architect Mr. Zack Ramjan to build a new system and remedy the shortcoming of compute power that was hindering its research. After assessing the needs and goals of the research groups within the institute, Mr. Ramjan decided that a hybrid HPC-private cloud solution based on the OpenStack framework would best serve his end users. Knowing that he and his team would be challenged to complete the project in the very ambitious six month timeframe that the institute had set, he looked to Silicon Mechanics to provide the skill and expertise necessary to get this key project completed within a six month timeline.

Solution Architecture


After accepting VAI's invitation to be the system integrator on the project, Silicon Mechanics' engineers began a thorough investigation into the institute's requirements and goals for both the short- and long-term.

Chief among their requirements was that the solution possess the accessibility and power to service the entire VAI community, which consisted of both researchers with technical computing skills and a great many novice users. The institute had a long history of utilizing HPC applications and platforms, meaning that most researchers were familiar with institute's legacy HPC platform and applications. These platforms had to be integrated into the new system in a way that didn't impact the user experience, while also introducing them to a new concept – HPC as a Service. At the same time, VAI wanted to provide its advanced users with the freedom and flexibility to experiment with system resources so they could tackle the complex research problems they faced.

Another major concern was overall solution management, requiring the ability to view resource usage metrics and flexibly adjust them as necessary. Because he oversaw the technology needs of over 30 different research groups, each carrying their own set of requirements, Mr. Ramjan needed to be able to simultaneously analyze VAI's cloud-based users and cluster-based users, identify the cloud and cluster user mix, and then dynamically tweak that number on the fly without massive downtime or retooling.

From a hardware point of view, expandability and scalability were also highly valued among VAI's requirements. This meant enabling a configuration that could support increased in workload size and quantity VAI is expecting in the future was a priority. It was also crucially important that solution offered strong GPU processing capabilities for researchers working with physics based equations.

With a thorough understanding of VAI's needs in place, Silicon Mechanics began working with its technical partner, Bright Computing, to architect a purpose built solution. Drawing on their own deep experience in the deployment of open cloud solutions, as well as by leveraging best practices promoted by the OpenStack community and Intel, the two partners were able to complete this phase and proceed to the solution build portion on schedule.





Build

With a clear roadmap in place, Silicon Mechanics began to build the solution and deploy it as designed. At the hardware foundation for the solution, Silicon Mechanics chose the Intel® Xeon® E5-2600 product family, due to its blend of efficiency, compute power, and SDI-centric features. To compliment this choice, a series of other Intel hardware components were chosen, including Intel enterprise SSDs and Intel network adapters.


To manage the hardware elements, Silicon Mechanics deployed the Bright Computing OpenStack system software. The Bright OpenStack platform provided the single pane of glass administration that VAI had been seeking, and due to its powerful ability to manage complex private cloud environments helped dramatically reduced the deployment time of the system. Deployment was so efficient that VAI researchers were able to start producing viable research data during the system's two week test period.

As the solutions integrator, Silicon Mechanics resolved a variety of technical challenges during the build and deployment phases, such as ensuring network traffic throughout the system was being properly supported and managed by an OpenStack Neutron server, and integrating Microsoft Active Directory into the system. Silicon Mechanics also supported the configuration of a highly-scalable HPC network attached storage (NAS) array provided by Data Direct Networks into the solution architecture. The strong working relationship between Silicon Mechanics and Bright Computing was key to the quick and efficient success of the project.

Result

Before final handoff of the solution to VAI, Silicon Mechanics conducted a set of rigorous validation tests to ensure that it operated to the technical requirements in the service contract. This testing included meeting both Silicon Mechanic's own proprietary testing standards and a set of testing guidelines set out by VAI. These guidelines included comprehensive functionality testing and ensuring integration of the HPC parallel file-system.

The testing process was completed in September of 2015, and immediately afterwards the system came online. As designed, the hybrid cloud HPC cluster has greatly simplified resource provisioning at the Van Andel Institute, striking the precise balance between power and usability that VAI wanted while minimizing downtime caused by managing, monitoring, and securing the system.



Hybrid HPC-Private Cloud Deployment Overview

Through the Bright dashboard, users are able to command system resources on-demand, placing them in a position of freedom and power. VAI users who are not technical computer experts, they have a painless means to provision and utilize the IT resources that they need to get their work done. For tech-savvy users, the solution has enabled them to construct their own operating system instance and run it in a cloud whenever they needed a specific type of computational resources. The lack of administrator involvement in the day-to-day usage of the solution has drastically increased efficiency at the institute. Mr. Ramjan had this to say, "Previously if someone wanted to add something new to the system, it had to be vetted, because any system change impacted every other user on the system. Cloud computing changes all of that. If one of our users wants to do something unique on the new system, who care? It's on the cloud. And we are fine with it."

Though the majority of the workloads – around 90% – are still run on the HPC cluster, there has been increased adoption of the cloud computing capabilities as users become more familiar with its potential. In some cases this new infrastructure has yielded great benefit, Mr. Ramjan added, "One big OpenStack success was getting a pediatric researcher off Amazon EC2 onto our hardware. He saw massive performance increases by going private-cloud thanks to our beefy hardware and fast network."


"Cloud computing changes all of that. If one of our users wants to do something unique on the new system, who cares? It's on the cloud.... And we're fine with it."

*-Zack Ramjan
High Performance Computing Solutions Architect, Van Andel Institute*

The overall reception to the solution has been extremely positive, both for the end users at the Van Andel Institute, who enjoy the power and flexibility it has brought, and from the administrators who see it as strategic and cost-reducing investment in the institute's sustained contribution to the field of genetic research. As VAI expands and develops in the future, the hybrid HPC-cloud computing system built by Silicon Mechanics is poised to grow with them. With its expansible architecture, anything from new compute nodes to additional storage can be easily brought online into the system with minimal burden to administrators, allowing the institute to dynamically respond to changes in scientific computing while also staying current with the trend toward cloud computing. Silicon Mechanics is very proud to have leveraged its strategic and technical expertise deploying cloud solutions to create this great success at VAI, and remains committed to supporting VAI's future growth as a trusted advisor.

"One big OpenStack success was getting a pediatric researcher off of Amazon EC2 and on to our hardware. He saw massive performance increases by going private-cloud thanks to our beefy hardware and fast network."

*-Zack Ramjan
High Performance Computing Solutions Architect, Van Andel Institute*



Silicon Mechanics, Inc., is a leading provider of servers, storage and high-performance computing technologies to the world's most innovative organizations. Since 2001 Silicon Mechanics has supported customers with its **Expert included.** approach, reflecting the company's passion for providing complete customer satisfaction and customer confidence in the return on their technology investments. Recognized as one of the fastest growing companies in the greater Seattle technology corridor, Silicon Mechanics is changing the way systems providers engage with customers.

To learn more, please visit www.siliconmechanics.com.

© 2016 Silicon Mechanics, Inc.

