Recreating the Power of a National Research Lab at a Smaller Scale – And Running it Without a Sys Admin

SILICON MECHANICS
SERVERS • STORAGE • HPC

By Andrew O'Neill, Engineer, Silicon Mechanics

Overview

One of the most interesting Silicon Mechanics™ projects I've been involved with in recent years was a very recent deployment. Earlier this year, we were asked to set up a smaller-scale high-performance computing (HPC) cluster for a researcher who had been a user of a much larger cluster Silicon Mechanics designed for a national laboratory.

Our task was to give him the same type of performance he was used to on the large cluster but on the smaller budget he was given for his specific project. This was an exciting design challenge, but it wasn't the only hurdle we had to overcome. We also had to keep in mind that he was given a budget for a robust system but not a system administrator – so he would be filling that role himself!

That meant the new cluster had to be much easier to run than the national lab cluster he was used to accessing as a user. And we also had to help him get all the system administrator training he needed to make his deployment successful. But our team was up to the task.

Problem

Several years back, Silicon Mechanics designed a large HPC Linux cluster for a U.S. government-sponsored national laboratory. Researchers at the lab used their cluster to conduct investigations into some of the most complex challenges facing the world today. To support those needs, their cluster had to have significant computing power, fast connections, and the ability to handle large data sets.

One researcher who used that large cluster was able to secure enough funding to build a new cluster dedicated entirely to his research project. He wanted to tailor the system to support his project's unique needs. The goal was to tailor the new cluster to his specific use case, changing the hardware slightly, so that he could get even better performance for his jobs than he was getting with the larger cluster.

The most unusual part of this situation, though, was that the researcher would now also have to become his own system administrator. That meant cluster management had to come to the forefront of this design. Even the most massive computing power means nothing if it's too hard to manage the hardware while also running your calculations.

And, of course, he needed the total project to fit within the fixed, much smaller budget he was given for his cluster and be up and running quickly so as not to fall behind on his investigation.

Solution

Our client found us by asking others at the lab who had designed the big HPC cluster he was used to working on. He reached out to us and shared his vision of a small but powerful cluster. Then we asked about specifics of the type of research he does and the types of datasets and calculations that involves, so that we could tune the system hardware to optimize his particular type of workload.

Tuning a cluster for a specific use case always makes it more effective (rather than making it so generic it can do some of everything fairly well, but nothing exceptionally well), so the Silicon Mechanics team was thrilled about optimizing the original design for him.

We took advantage of our deep partner relationships (NVIDIA® Elite Partner status, etc.) and were able to maximize the client's budget while including all the components he needed to achieve his goals. And, because we have in-house, U.S.-based production facilities, we were able to quickly create customized servers that were optimized for the results we were looking for.

System Hardware Design

We designed and built a six-node Linux cluster powered by refreshed Intel® Xeon® Silver second generation central processing units (CPUs) working in parallel with NVIDIA V100 Tensor Core graphics processing units (GPUs). These GPUs were the most advanced available at the time. They are capable of breaking the 100 teraFLOPS (TFLOPS) barrier of deep learning performance and offering the performance of up to 32 CPUs in a single GPU so this was a great choice for the high performance our client was looking for.

The team also designed and built a Lustre file system-based data storage array. The parallel, distributed Lustre file system is used by a good chunk of the top ten supercomputers and in many large-scale clusters. They like Lustre for its scalability but Lustre also offers very good input/output (I/O) performance, which is critical for researchers with large datasets like our client. It also has high availability features, such as robust failover and recovery mechanisms that are project lifesavers when things go wrong with the power or a hardware fault occurs.

The network was based on 200 gigabit per second (Gb/s) Mellanox® high data rate (HDR) InfiniBand Quantum™ CS8700 switches operating on 200 gigabit HDR cables. InfiniBand provides end-to-end throughput over 25 GB/s across 200 Gb interfaces empowering multi-node compute with very high throughput and very low latency, ideal for large data sets.

Cluster Management

The biggest challenge for the Silicon Mechanics team, frankly, was what to do about the unusual position the client was in, with regard to becoming a system administrator for the first time. Our first step was to suggest the cluster be provisioned with the Bright Cluster Manager from Bright Computing.

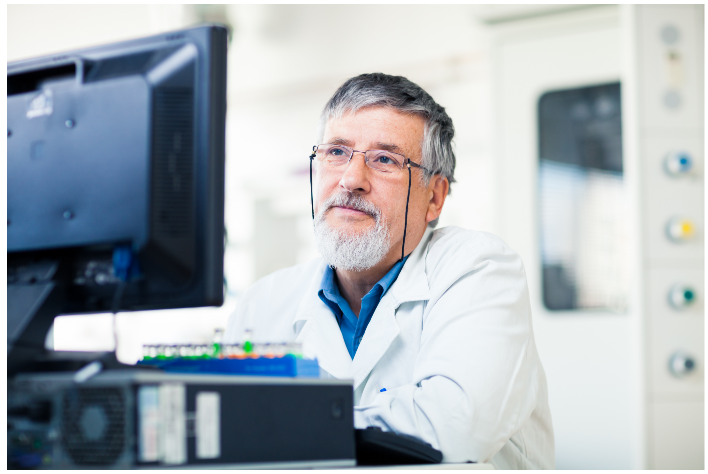

The client had heard of this enterprise-grade software system before but had not used it very much. We discussed it with him in depth and, when he was comfortable with the idea, deployed the software. Then, to meet his specific research needs, we also installed the Bright Computing Data Sciences add-on.

The Bright Cluster Manager offers a web-based graphical user interface in addition to the more traditional, but more technically complex, command line interface (CLI) used by system administrators. The intuitive software would make it easy for our client to understand how to provision nodes, install and update applications, setup and maintain separate node images and queues, which also helped simplify a complex Lustre deployment.

Becoming a System Admin

Beyond the installation of the project cluster itself, the Silicon Mechanics team had mapped out a hands-on training program for our new system administrator. However, that plan became impossible because of travel restrictions related to the onset of COVID-19. Instead, we suggested a rigorous online-only training program. The client agreed and the engineering team quickly created an online-only curriculum that replicated the situations a systems administrator would encounter on this specific system.

Online-only training can be less than ideal or even downright challenging in some situations. However, the customer embraced the online curriculum and ended up being very successful in mastering the training. It helped that the Silicon Mechanics engineering group that setup and configured the cluster prior to shipment also designed and conducted our online training. The team lead made sure that the training not only reflected the hardware design but also the needs of a new system administrator.

Results

Our “zero defect” build process and standards require technicians to hand inspect every order at multiple points during the build and testing process. Plus, each system is subjected to rigorous additional testing by a proprietary suite of automation and stress testing tools. This means systems arrive 100% operational, reducing time-to-result. In this case, because of COVID-19 travel restrictions, the fact that the system was ready to go on day one was not just desirable but critical.

Now, the client is happily using the new private cluster rather than waiting in line for the lab’s larger, shared system. And he is putting his new system administrator skills to good use managing the cluster software as it does data analysis across various areas of research.

Silicon Mechanics provides robust product documentation for everything we work on. The goal with documentation is to allow customers to easily replicate hardware settings, firmware updates, and software versioning. Ensuring standardization when installing new systems leads to efficient system deployment. In cases like this, though, the value of putting in that kind of documentation effort really paid off. (Though the client was glad to hear we also have expert U.S.-based service level agreement (SLA)-based customer support, in case something comes up that he can’t manage on his own.)

This unusual deployment went so well, in fact, that the internal team here at Silicon Mechanics thinks we can use our enhanced remote learning capabilities for more clients, not just while COVID-19 causes social distancing but when budgets or other reasons prevent on-site training. That should really expand the number of customers who can take advantage of the power of HPC systems, which is a great thing for everyone.

Solution

Silicon Mechanics, Inc. is one of the world’s largest private providers of high-performance computing (HPC), artificial intelligence (AI), and enterprise storage solutions. Since 2001, Silicon Mechanics’ clients have relied on its custom-tailored open-source systems and professional services expertise to overcome the world’s most complex computing challenges. With thousands of customers across the aerospace and defense, education/research, financial services, government, life sciences/healthcare, and oil and gas sectors, Silicon Mechanics solutions always come with “Expert Included”SM. Learn more at www.siliconmechanics.com.