Unique turnkey appliance offers privacy, security, version control, and file transfer speed

When Knome needed a turnkey hardware platform for its knoSOFT™ genome interpretation application and kGAP™ informatics engine, it turned to Seattle-based solutions provider Silicon Mechanics. Knome had been using Amazon Web Services (AWS) for computationally intensive tasks, but needed to provide a locally-installed system with full control for clinically-oriented customers, who are sensitive about security, version control, and file transfer times.

Working closely with Silicon Mechanics, Knome developed the knoSYS™100, which is a combination of Knome's software and a hardware platform configured by Silicon Mechanics. The system is designed to provide labs with an end-to-end solution for the interpretation of human genome, exome, and targeted sequence data.

Seeking in-house solution

Image 1: Dr. George Church, co-founder of Knome, with the knoSYS™100

Cambridge, MA-based Knome provides human genome interpretation systems and services, working with pharmaceutical companies and medical researchers to understand the genetic basis of disease, tumor growth, and drug response. In its Interpretation Services division, the company has worked with researchers on a wide range of diseases, including Parkinson's, Alzheimer's, asthma, and cancer. Typically, the informatics for large interpretation projects, covering hundreds to thousands of genomes, would be processed in the cloud using AWS.

In 2012, as the market began to turn toward clinical applications, Knome recognized the need for labs to have internal interpretation capabilities. Instead of interpreting a large number of genomes in order to understand a specific disease, the requirement is to address the medical needs of one individual or family. In the clinic, processing and comparing large numbers of genomes is not as important as turnaround time, privacy, security, and version control.

To meet the demands of the newly emerging clinical market, Knome decided to develop an end-to-end system that would enable a lab to effectively handle the computational and interpretation requirements of next-generation sequencing-based tests within their own facility.

One of the most basic issues to be tackled is the transfer of files that are often over a terabyte in size. To transfer files this size to the cloud is typically too time consuming to be practical in a commercial environment. Even a relatively small file of 200 gigabytes can take two days to upload to the cloud. In addition, clinics have significant concerns regarding security and privacy; most early adopters want to keep their patients' genomic data behind their firewall.

"Ultimately, privacy, security, and transfer speed drove us towards an in-house approach, so we began the process of deciding what hardware would be needed to perform all the computational tasks that AWS did for us in the cloud," said Michael McManus, Senior Vice President of Knome.

Searching for the right hardware

Configuring the hardware required to optimize the intensive computational tasks required by whole genome-level informatics was no easy task. Knome's knoSOFT™ and kGAP™ software packages are not single applications—rather, McManus likens them to an "ecosystem" of about 45 applications all running together.

To design, supply, and support the knoSYS™100's hardware platform, Knome turned to Silicon Mechanics, a provider of rackmount servers, storage, and high-performance computing solutions. Silicon Mechanics collaborated with Intel on product development activities to create the ideal solution. As a premier reseller of Intel products, Silicon Mechanics receives an advanced roadmap of Intel technology, as well as engineering samples of hardware to test for product compatibility, enabling it to base its product lines on Intel's planned product launches and end-of-life plans.

Knome had already selected Intel processors, primarily because of their fast RAM throughput and the availability of hyperthreading, which improves parallel processing across compute cores within a processor. In addition, Intel worked closely with Knome to help them choose the right chipset for the appliance. The team reviewed the breadth of options in several of Intel's chipsets and looked at the various performance bands with different characteristics.

According to Knome's McManus, the evaluation began with the basics, like processor speed, the I/O capabilities of the hard disks in the compute nodes, and the database connections of the nodes.

"There were many discussions on bits and bytes and how to ensure we get the performance we need, as well as details like the best way to organize hard drives in a storage array to minimize failure rates and maximize performance," stated McManus. "The Silicon Mechanics team was extremely helpful in openly working through the various issues we encountered."

Silicon Mechanics began by developing a deep understanding of Knome's requirements so it could narrow down all the available processor options in order to optimize price and performance. Power consumption, temperature, noise, and application performance were all evaluated.

"We went through a lot of discussion and used a process of elimination to narrow down all the available processor options to the optimal model for Knome's users," said Tim Groen, Silicon Mechanics' Senior Enterprise Account Manager. "We started by asking about requirements in a pretty broad way and prepared a comparison of twenty possibilities in terms of power, price, lifecycle, and availability."

Image 2: knoSYS™100 integrates seven servers in a rack, all optimized to run Knome's interpretation software and informatics engine.

For this type of custom-designed product, Silicon Mechanics is often called upon to bridge the gap between major vendors, component suppliers, and end users. Silicon Mechanics' knowledge of Intel's product lines was essential for educating Knome on the variables that were important for the application.

Hardware platform gives end users I/O support, computational power, and fast file transfer at an affordable price

The hardware ultimately selected for Knome's new system, the knoSYS™100, is a sophisticated high-performance grid computing system that integrates seven servers in a rack—all optimized to support 10 or more simultaneous users running Knome's interpretation software and informatics engine. One of the key features of the appliance is that it uses industry-standard components, including the Intel® Xeon® Processor E5-2600 product family. This approach greatly reduces upfront capital and long-term support costs as compared to proprietary hardware solutions.

Within the Intel Xeon Processor E5-2600 product family were options with six- or eight-core processors, and Silicon Mechanics ultimately selected an option whose features offer the maximum RAM throughput: a high-performance computing cluster with four nodes, each with two 8-core/16-thread, 2.4 GHz, 64-bit Intel Xeon Processors E5-2665 with 20MB cache, capable of 500

GFLOPS per node. The Intel Xeon Processor E5-2600 product family offers performance gains up to 80 percent over a previous generation Intel® Xeon® processor-based server and provides the reduced I/O latency Knome required with Intel® Integrated I/O.

Storage is provided by a Lustre® array. According to Knome's McManus, Lustre is being used because its I/O rate is matched to the systems' intensive computational capacity. "We have all these applications running—some need memory, some need fast disk interaction, some just need a lot of CPU," says McManus. "Each application has its own characteristics that must peacefully coexist with the others. Lustre gives us enough I/O support to manage a wide variety of usage scenarios."

He is also pleased with the appliance's hefty 512 gigabytes of RAM and the eight powerful Intel Xeon processors. "In this box format, we were able to create an architecture equivalent to eight of the large machines we use on AWS, giving us significant computational power at a much lower lifetime cost."

"With this solution, we have something that is fast, reliable, easily supported—but not super expensive," said McManus. "We have worked with a number of high-end custom hardware vendors, which may be a good solution for some customers, but the price is far higher. We decided we could get the best bang for the buck with a flexible, industry-standard solution like that offered by Silicon Mechanics."

Silicon Mechanics integrates and tests each individual server; installs the operating system, cluster middleware, and the Lustre array; racks the servers; then integrates the power, network access, switching, and cables.

Knome is looking forward to the day when precision medicine will considered standard practice and the broad use of genomic information will be routine. Systems like the new knoSYS™100 are poised to play an important part in this revolution.

About Silicon Mechanics

Silicon Mechanics, Inc. is an industry-leading provider of rackmount server, storage, and high-performance computing solutions. Deploying the latest innovations in hardware and software technology, we work in collaboration with our customers to design and build the most efficient, cost-effective technology solution for their needs. Our guiding principle, "Expert included," is our promise that reflects our passion for complete customer satisfaction, from server and component selection to superior installation and ongoing technical support.

Silicon Mechanics has been recognized as one of the fastest growing companies in the Greater Seattle Technology Corridor.

About Knome

Knome Inc. is a leading provider of human genome interpretation systems and services. We help clients in two dozen countries identify the genetic basis of disease, tumor growth, and drug response. Designed to accelerate and industrialize the process of interpreting whole genomes, Knome's big data technologies are helping to pave the healthcare industry's transition to molecular-based, precision medicine.