

Van Andel Institute Integrates Traditional Grid and Cluster Computing with Cloud Computing using Bright Cluster Manager

CASE STUDY

“We know that cloud computing is the wave of the future. The hybrid approach we are getting with Bright is providing a path that helps us transition. ”

— Zack Ramjan
Research Computing Architect at VAI

Van Andel Institute (VAI) is an independent biomedical research and science education organization based in Grand Rapids, Michigan. VAI hosts thirty individual research groups who use genomic sequencing analysis, molecular dynamics simulation, and modeling to investigate epigenetics, cancer, and neurodegenerative diseases. Recently, VAI deployed a unique new hybrid infrastructure featuring Bright OpenStack systems management software.

The Bright OpenStack deployment wrapper lets VAI manage both high-performance compute (HPC) grid and cluster computing and cloud computing within the same infrastructure, greatly reducing the labor and effort needed for management and change control. Perhaps more importantly, it also helps VAI respond dynamically to the accelerating trend toward cloud computing they see coming down the highway.

“Grid and cluster computing has been the standard for years, but we know that cloud computing is the wave of the future,” said Zack Ramjan, research computing architect at VAI. “The hybrid approach we are getting with Bright is providing a path that helps us transition from one to the other.”

The Challenge

Zach Ramjan considers his main challenge finding a way to meld hardware, software, and storage into a system that can handle the massive amount of data generated by VAI’s varied research groups.

To solve that challenge, he sought an environment that could handle the massively parallel processing (MPP) and analysis produced. “That’s where Bright comes in,” says Ramjan. “We have the cluster and grid and will eventually have the cloud, so we can use a “divide and conquer” approach, efficiently assigning tasks among 50 computers and the cloud resources.”

Another challenge he faces is that the thirty different research groups are working on highly varied projects, each with unique requirements. Some of the researchers are doing genomics and others are doing simulation – and they needed a solution that would work for all of them.

Ramjan explains that access to a cloud approach puts the user in the driver seat, so he would not have to develop a single solution that makes everyone happy. Each user can have their own solution, carving off their own virtual piece of the pie. “Bright is helping us manage that variety of approaches. With both the legacy cluster mode and the cloud mode, we are creating an environment in which it is easy for users to come on board and do their specific work efficiently.”

The Solution

The total solution includes three key components: Bright OpenStack software; 43 compute nodes, representing 1100 CPU cores, provided by Silicon Mechanics; along with storage and data hosting supplied by Data Direct Network. Bright OpenStack is the brains of the system and integrates the hardware, managing the CPU units and the storage devices. The user interacts with the Bright dashboard, and Bright OpenStack interacts with the physical elements of the system.

An expert with many years of deploying HPC systems by hand, Ramjan decided to opt for the Bright solution because he did not have the time or the manpower in house to quickly set up a hybrid HPC-cloud system. He reasoned that engineering its solution would take years. "We are actually kind of lucky that we started from nothing because we had no legacy baggage, so we could design the solution as we saw fit from the beginning, and we did that with the help of Bright and others."

Bright partner and system Integrator for the project, Silicon Mechanics, was confident the Bright tool set would be a successful way to tackle the challenge. "As we worked with VAI to define their system architecture, we looked to the Bright OpenStack system management software, knowing its strength in managing complex, private cloud computing environments," said Daniel Chow, COO/CTO at Silicon Mechanics. "We are excited to see the joint solution empower VAI researchers to accelerate their scientific findings."

A key feature of the Bright OpenStack software is the ease of management and change control. Ramjan notes that it is very difficult and time consuming to scale up by managing many machines

one by one. "With Bright you can see the entire resource, or see your hardware through a single pane of glass. Right now our HPC to cloud ratio may be 90-10, but tomorrow we know our end users are going to be more cloud-centric. With this solution, I can dynamically pull resources into the cloud portion, but if the next day it turns out there's less cloud demand, we can pull it back. We can dynamically shift that ratio as we see fit without any down time."

According to Ramjan, several other institutions he has spoken to also have their eyes on transitioning to an OpenStack cloud. However, those larger institutions with significant investments in "old school" HPC infrastructure say they are probably three to five years away from doing so. "Because they have such a complex environment, everything has been custom designed in-house and changes must be done with their own labor, which makes them less flexible. With Bright, we bought the solution and everything came with the package. It definitely put us way ahead of the game."

The Result

The HPC cluster went online in September 2015 and is already highly utilized, which Ramjan considers a good sign. Bright makes it easy for other team members to jump in to provide assistance to users without needing his intervention. He says that Linux can be quirky technically, so without Bright OpenStack, even basic tasks would have fallen solely in his lap, which would have slowed down deployment considerably.

"We know that scientific workloads do not get smaller every year, but are constantly expanding. From our experience, the size of data continues to grow exponentially. We have more

than 40 compute nodes representing 1100 CPU cores today – but what about next year when we get to 2,000 or 3,000 cores? We wanted an expandable and scalable solution – this is a core capability of Bright. The management tool makes it easy to buy new equipment, take it out of the box, put it on the shelf, and plug it in. Bright can pull it right into the existing environment."

The cloud portion has recently gone live and is available to VAI's users. Ramjan says that VAI's early adopters, typically power-users who are quite savvy, are starting to appreciate its value. He expects use of the cloud to grow in popularity as others get more familiar with the resource.

Although VAI's existing workloads are 90 percent grid and cluster, they expect to move toward the cloud in the future. Bright is giving them an expandable and scalable turnkey solution that lets them combine HPC workloads with big data analytics workloads in the same infrastructure, and to have the choice of working in either a bare metal or virtualized infrastructure. It's also providing a path that helps them transition from one to the other.

Concludes Ramjan, "We were very, very pleased with Bright's expertise and their grasp of the situation. You will find plenty of people who want to sell you some sort of cloud solution that they probably make right after they take your order, and hope that it works out. But with Bright, it's a lasting relationship. They've taken really good care of us, they know what they're doing, and we are extremely satisfied."