BIOS IT

CASE STUDY: Montana State University

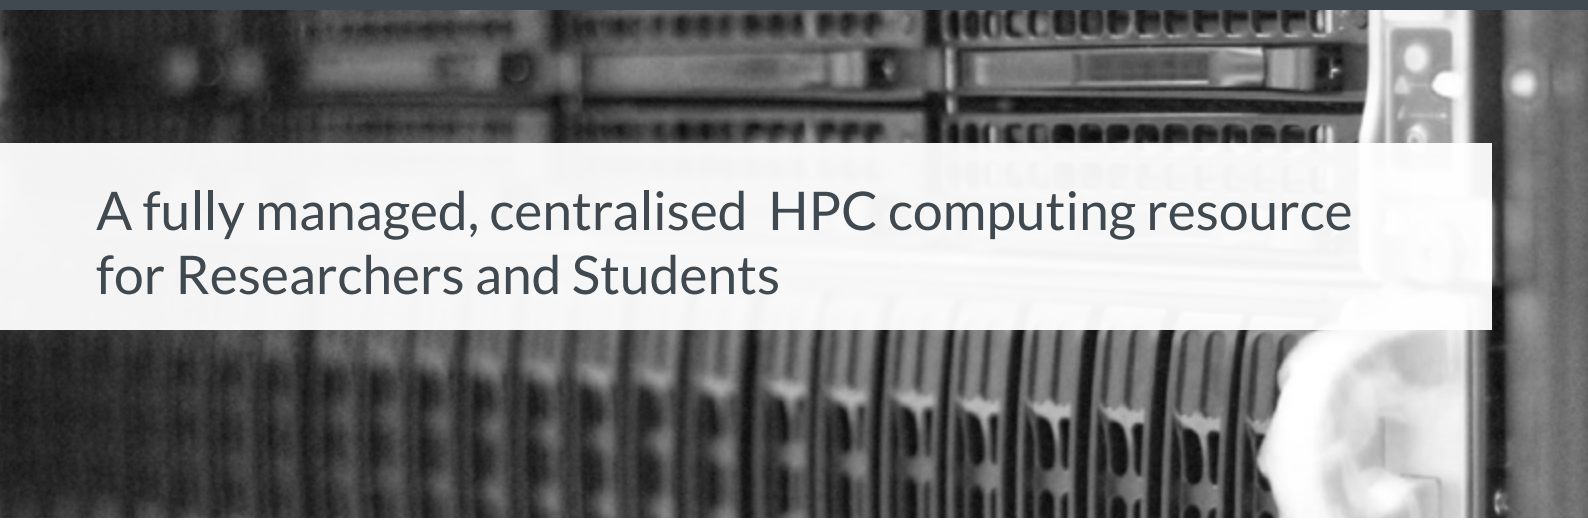

A fully managed, centralised HPC computing resource for Researchers and Students

BIOS IT delivered a centralized HPC cluster for Montana State University to conduct their research from a central compute and data repository. The system features a highly-available high performance Lustre storage solution which facilitates vast quantities of data to be processed and analyzed. BIOS-IT took advantage of virtualization and cloud computing technologies to provide a flexible and affordable solution for the diverse needs of their researchers. As part of an on-going HPC managed services agreement with Montana State University, BIOS-IT provide the day to day administration of the cluster and remote configuration assistance, allowing the HPC staff at MSU to focus on enabling users and broadening the adoption of HPC among the scientific community.

BACKGROUND

MSU is recognized nationally for its research prominence among leading public research universities. Its prolificacy has led to many significant and even world-changing discoveries. MSU holds more than 250 active technology licenses. In addition, 91 patents and 36 plant variety certificates have been issued for MSU discoveries with many more pending.

Research expenditures at MSU typically exceed $100 million annually. Particularly notable research departments include immunology; chemistry and biochemistry; transportation; physics; and land resources and environmental sciences. Much of the research funding is derived from the National Science Foundation, National Institutes of Health, the Departments of Energy, Defense and Agriculture, and private sources.

THE CHALLENGE

MSU's Information Technology Center's Research Cyberinfrastructure(RCi) group was tasked with building a shared community computing resource. Their goal was to build a flexible infrastructure that would let them offer compute and storage services at an affordable

MONTANA STATE UNIVERSITY

AT A GLANCE

15600 STUDENTS
AND GROWING

$100 MILLION
ANNUAL RESEARCH EXPENDITURE

OVER **250** ACTIVE
TECHNOLOGY LICENSES

91 PATENTS ISSUED
FOR MSU DISCOVERIES

"BIOS IT delivered a centralised HPC cluster for the whole University. The UI has been great. Our users submit and monitor their own jobs with ease, and we can focus on monitoring and managing the cluster with increased facility and insight."

Pol Llovet,
Associate Director of
Cyberinfrastructure

price. It had to be scalable in small increments, easy to administer, and able to use off-premise resources to extend the infrastructure when necessary. It also had to be easy enough for students to use, and reliable enough for researchers to count on for their grant projects.

THE SOLUTION

The team at MSU spent a considerable amount of time researching a solution. Their approach to providing infrastructure is an interesting one. Rather than build a giant cluster and house it in a data center on campus, MSU chose to build a modest cluster of generalized nodes on site, and take advantage of technologies such as virtualization and cloud computing to provide a fl exible, aff ordable solution that meets the unique needs of their diverse researchers. Since these nodes would have to be reconfi gured quickly and reliably to meet the needs of various projects, MSU required management software that could handle the task. They chose a BIOS-IT cluster solution powered by Supermicro, Mellanox, Intel® Enterprise Edition for Lustre, and Bright Cluster Manager. The cluster provided the university with a central HPC capability for a diverse range of workloads with high performance fabric for parallel workloads and a 500TB parallel fi le system to handle large IO demand and complex data analysis.

BIOS-IT selected Bright's management software as it makes it easy to confi gure the overall cluster initially, and reconfigure nodes as needed. The cluster was pre-installed at the BIOS-IT HPC facility in New Jersey and soak tested for a number of days, using BIOS IT's own suite of software tools. This reduced the installation time on site and meant the system was ready for use from the researchers and end users almost from day one.

FUTURE PROOFING WITH OPENSTACK

MSU's blueprint called for a modest-scale on-site cluster, virtualized clusters, and the ability to expand use of off -premise resources when necessary. vScaler was the solution of choice. It enables MSU to spin up virtualized HPC clusters on existing servers, and easily expanding to managed cloud servers. Working with MSU, BIOS-IT selected vScaler OpenStack for a number of reasons: vScaler is a certifi ed OpenStack Powered Platform, ensuring interoperability; it was designed to make spinning up HPC clusters within an OpenStack cloud environment easy. It also enables MSU to manage their local HPC cluster, their remote HPC clusters and the OpenStack cloud itself from the same powerful user interface.

THE RESULT

For MSU, BIOS-IT's implementation provided a seamless solution. It off ered the university a customized & all-inclusive software for deploying, provisioning, monitoring, health checking, and much more. It was easy for sys-admins to use, and it was easy for end users to access.

In collaboration with MSU, BIOS IT have implemented their Remote Cluster Administration (RCA) managed services package to help MSU continuously monitor, optimise, health check and support their cluster. This partnership has helped to stream line the Research Computing Team's workload and allowed them to focus their eff orts on advertising the benefi ts of the centralised facility to the Researchers and ultimately accelerating the advancement in their discoveries.

The installation went according to plan. It was completed in January 2015, and has been running user jobs and producing results ever since. There are roughly 20 labs using the centralized HPC facility, including mechanical engineering, physics, neuroscience, ecology, computational chemistry, and others. More users are being added every month.

The new system is meeting the goals and expectations of MSU's team. Labs are able to access the compute resources they need without the cost and overhead of setting up their own clusters. And the centralized nature of the cyberinfrastructure has granted more faculty and students access to compute resources than ever previously envisioned.