CASE STUDY

High Performance Computing (HPC)
Intel® Xeon® Scalable Processors
Intel® Omni-Path Architecture

Shandong University

Deploying High Performance Computing for Environment as a Service to Support a Diverse Computing Audience

Shandong University supercomputer at a glance

- First HPC environment in China to offer service in the Cloud, running multiple traditional HPC and non-traditional HPC workloads with Intel® Xeon® Scalable processors and Intel® Omni-Path Architecture (Intel® OPA)
- 380 teraFLOPS of double-precision floating-point performance (e)[1] with 1.6 PB of storage
- New supercomputing center meets the requirement of scientific research needs in HPC, Cloud, Big Data, Artificial Intelligence/Deep Learning, and Data Analytics
- An HPC cloud platform with hybrid architecture, containers, and mobile app technology

Executive Summary

Shandong University, founded in 1901, is one of the oldest and most prestigious universities in China. It is the second national university established in the country and one of the first in China to install High Performance Computing (HPC) resources. The School hosts the Shandong Center for High Performance Computing, an HPC and resource sharing platform established in 2002. It provides an environment for world-class modern research for fundamental science, material science, bioscience, environmental science, and computing, including grid technology, parallel computing, mass data processing, cryptanalysis, and virtual reality and visualization technology. The center is a milestone for the national computing environment and a critical component of the ChinaGrid project, one of the world's largest grid computing implementations.

Challenge

HPC resources in Shandong University are needed across a diversity of learning disciplines and environments, and to support national initiatives. The insights needed to support China's ongoing 5-year plans have leveraged HPC resources. The Shandong Center for High Performance Computing has undertaken some key research and development programs under the Eleventh, Twelfth and Thirteenth Five-Year Plans. It is also part of the National 863 Plan, a program established in 1986 to stimulate technology development in China.

The supercomputing center supports research across Artificial Intelligence and Machine Learning (AI/ML), experimental teaching and virtual/augmented reality, big data and others, serving both sophisticated and unexperienced users. Thus, Shandong University recognized the need to provide computing resources that extend beyond traditional simulation and modeling used by the empirical sciences. To meet the needs of a hugely diverse user audience, the center focused on building their next HPC system to provide Environments as a Service (EaaS).

Running as EaaS, the new supercomputer needed to support multiple operating systems (OS), various software versions (not just the latest one), deep learning frameworks, and more that could run on the x86 instruction set processors and GPUs. The hardware and software needed to be easy to manage and operate for both system administrators and users. The solution had to provide both large-scale and small-scale HPC cluster computing and powerful desktop-like environments—all enabled through user-focused interfaces that simplified and accelerated each environment deployment.

Solution

In designing their HPC system, the Shandong Center for High Performance Computing employed smart microcode and container and mobile application technologies on a cloud service platform all based on a hybrid architecture. To support a sophisticated environment that was user-friendly yet able to support

Shandong University's new system incorporates Intel® Xeon® Scalable processors interconnected by Intel® Omni-Path Architecture fabric.

a wide base of research needs, open sharing, and efficient management, their software included bar code scanning. The enhancements will simplify user logins, enable social-based mobile applications to push notifications to users, and provide an environment that allows self-administration of systems, environments, applications, and data for each user.

The project began in March 2017. Built by Huawei and Clustertech, the new system includes 172 nodes of dual-socket Intel® Xeon® Gold 6132 processor interconnected by Intel® Omni-Path Architecture (Intel® OPA) fabric. The cloud service platform delivers 380 teraFLOPS of performance (e)[1] with 1.6PB storage capacity. It was jointly launched in July 2018 by Huawei, Clustertech, Intel, and the university.

System Management software provides one-click configuration and installation and batch installation, and supports dynamic capacity expansion or reduction based on the service traffic. It's also provides intelligent power consumption management. It can monitor, and analyze, and diagnose various energy efficiency indicators, and take action based on the analysis and diagnosis results to reduce power consumption. The software also supports centralized monitoring and unified management of various devices.

Per Huawei, the infrastructure provides board-level to system-level energy-saving measures, intuitive real-time monitoring, and dynamic energy-saving technologies to reduce power consumption by up to 40 percent[2]. The system-level energy-saving measures include:

- Efficient uninterruptible power systems (UPSs)
- In-row air conditioners
- Frequency-conversion cooling
- Modular design
- Natural cooling
- NetEco intelligent power consumption management software

These measures decrease the overall power usage effectiveness (PUE) to less than 1.2.

Results

Since deployment, the new system has supported projects running a wide range of OSs, parallel workloads, AI/ML jobs, data analytics, and more.

The new system leverages widespread use of mobile devices by integrating mobile services for authentication, self-administration of users' workloads and data, and push-notifications of job activities and status. This allows users to have greater awareness and control of their projects running on the new system.

Meeting the needs of a very wide user base across multiple research areas and computational applications, the system is built for a wide variety of workloads. TensorFlow* and Jupyter are installed for deep learning and AI applications; several bioinformatics tools support easy biodata analysis

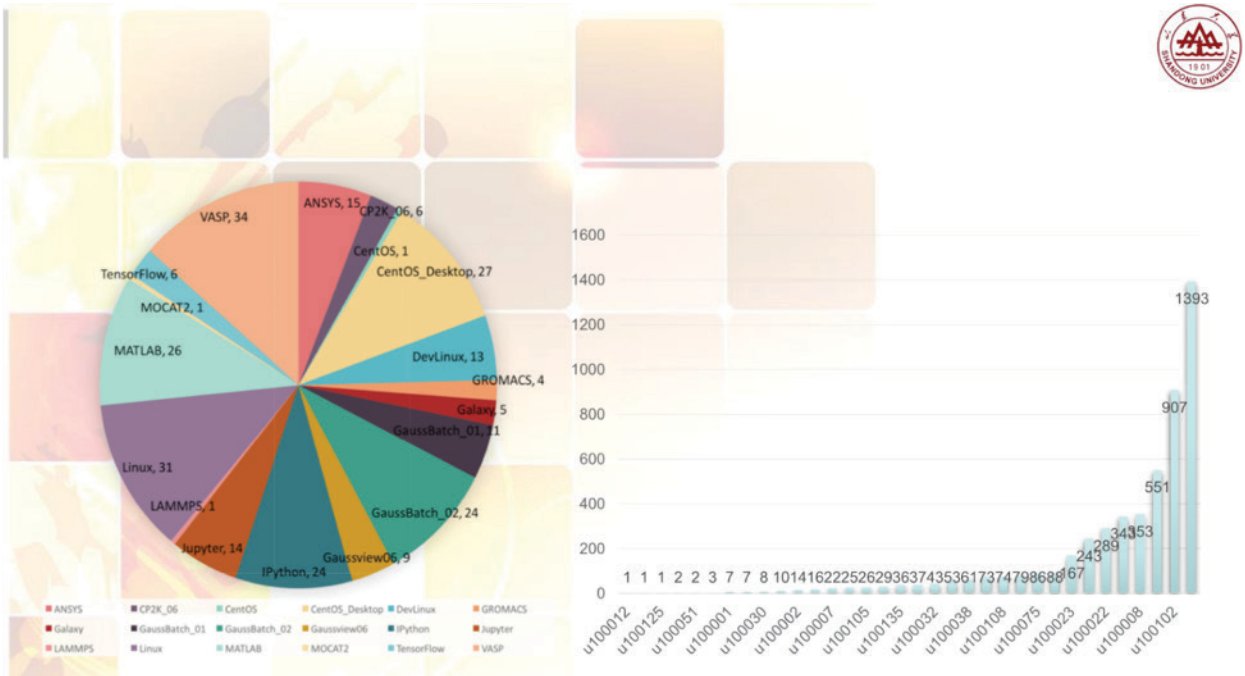

Figure 1. Recent environments and workloads

workflows. The cluster has become a public open platform that integrates various biological information analysis functions, such as data uploading and processing, sequence alignment assembly, sequence analysis, SNP/WGA analysis, and data visualization for bioinformatics.

The new cluster also supports traditional computational sciences, including computational chemistry with applications like Gaussian and GaussView, enabling building, analysis, and visualization of complex molecules and materials. And, supporting the ChinaGrid distributed computing model, users can request cluster resources that the system then orchestrates into virtual HPC clusters for their jobs, all through a sophisticated yet easy to use queue management system.

Solution Summary

Shandong University's Center for High Performance Computing needed their next HPC resource to serve a wide diversity of users with a range of computer experience and computing needs. They deployed a 172-node cluster running a sophisticated stack of software to support traditional HPC jobs, modern research in AI/ML, analytics, and bioinformatics, and non-traditional workloads and personal desktops in an Environment as a Service model. The cluster was built on Intel Xeon Gold processors and an Intel OPA fabric.

Where to Get More Information

Learn more about Shandong University at http://www.en.sdu.edu.cn/.

Learn more about Intel Xeon Scalable Processor and Intel Omni-Path Architecture at https://www.intel.com/xeon and http://www.intel.com/fabrics.

Solution Ingredients

- Intel® Xeon® 6132 Gold Processors
- Intel® OPA fabric
- Server: Huawei FusionServer 2488H V5/ Huawei FusionServer 1288H V5 172
- Storage: Huawei OceanStor 2600 V3
- Filesystem: Lustre*
- System Management: Huawei eSight
- Infrastructure: Huawei Fusion Module 2000

[1] Note that "e" means "estimated"; Performance measurement comes from the calculated theoretical Linpack performance based on the CPU and nodes number. HPL Linpack Rpeak is: 2.6GHz*14*2*32*172=400TFlops, over 380TeraFlops. System configuration: Huawei FusionServer 1288H V5/ Huawei FusionServer 2488H V5 *172 with Intel Xeon 6132 Gold Processors (14Cores/2.6G/140w), Intel OPA fabric, Huawei OceanStor 2600 V3 *2 (8*80TB HDD) and related 300TB system disk, Lustre, Huawei eSight, and Huawei Fusion Module 2000.

[2] In Huawei Fusion Module 2000 system, board-level liquid cooling PUE is about 1.1, and the average air-cooled PUE is about 1.6, so heat dissipation efficiency is improved by about 40% [(1.6-1.1)/1.1]. Source: Huawei

0818/SY/RJMJ

Please Recycle

337633-002US