CASE STUDY

High Performance Computing (HPC)
with Intel® Omni-Path Architecture

Penguin Redefines HPC in the Cloud

Intel® Omni-Path Architecture with Intel® Xeon® Processors deliver exceptional bare-metal performance for HPC cloud services with Penguin Computing On Demand (POD)

Penguin On Demand

- Bare-metal HPC clusters with outstanding HPC performance
- Built on Intel® Xeon® Processors and Intel® Omni-Path Architecture fabric
- Predictive performance and costs for HPC on demand in the cloud
- Highly experienced, on-site HPC experts support customers for HPC coding optimizations

Challenge

Eight years ago, Penguin Computing saw the growth of cloud computing. It was clear that people were trying to use the cloud for HPC. "But the clusters they built did not perform well, because cloud networks are largely virtualized," said Phil Pokorny, Chief Technology Officer at Penguin Computing. "It wasn't possible to get a good guarantee of locality or bandwidth between instances with such an infrastructure." Today, according to Pokorny, building an HPC cluster in a traditional cloud environment, even with a 10 GbE network, does not scale well to the number of nodes that the typical HPC customer needs to run big projects.

Besides the performance and scalability, Penguin found that HPC in the cloud was expensive. Service providers charge for data going in and coming out, and for clusters that were configured but not busy. Additionally, it wasn't cost effective to tear down a non-busy cluster and stand it up again later. Finally, in the cloud, customers were on their own when it came to HPC. For the relatively new user or user who was more scientist than programmer, a lot of time could be spent just ramping up, and that meant costs not running jobs. HPC in a traditional, virtualized cloud environment wasn't working—at least not like an on-premise deployment, which experienced HPC users—even non-programmers—expected.

Solution

Penguin Computing has been providing HPC clusters for years to many customers. Today they are building some of the largest, fastest, and most sophisticated systems in the world—from the first Commodity Technology Systems (CTS-1) at the Tri-Labs (Lawrence Livermore National Laboratory, Los Alamos National Laboratory, and Sandia National Laboratory) to clusters at commercial enterprises in oil and gas, life sciences, and engineering, to name a few. "When we looked at HPC in the cloud," commented Victor Gregorio, SVP of Cloud Services at Penguin Computing, "we decided that, with our partner solution offerings, professional services, and general breadth of expertise in the industry, we were in an ideal position to create an on-demand cloud environment for HPC workloads." Since its launch over eight years ago, Penguin Computing on Demand (POD) has grown in customer use and demand, driving expansion of the clusters in the POD service.

A Multi-Level Service Approach

Penguin's solution is designed around four important requirements: a bare-metal infrastructure, easy access to a wide range of software, transparent billing, and professional HPC support.

Bare-Metal Performance—Critical to their approach from the beginning was to present an HPC cloud offering that ran on bare metal—not virtualized. To maximize performance, nothing within the compute domain is virtualized. Penguin runs the service just as it would be run on-premise at any laboratory or in any enterprise IT

data center. "When a user signs up at pod.penguincomputing.com, logs in, and submits a job, it is scheduled on a partitioned number of nodes, just as if the cluster was in-house," added Gregorio.

Range of Software—To meet the needs of a broad range of engineering and scientific domains, Penguin also provides a range of commercial and open source HPC applications and codes, from computational fluid dynamics (CFD) to vehicle crash analysis to weather prediction, and more. "We have worked with the commercial ISVs to create a cloud licensing structure that makes it easy for customers to run their jobs on POD clusters," stated Pokorny. "Customers can even license their own software to a POD cluster for a particular job."

Predictable Billing—Billing is very transparent and predictable. Penguin charges by the core-hour, without adding costs for additional tasks, such as uploads and downloads. "If a customer knows how many cores they'll need and can budget how long their job will take, they can easily estimate the charges," added Gregorio.

Professional Support—Not all HPC users are parallel programmers. So, Penguin provides deep expertise on optimizing codes for the clusters. Each member of the POD support team has five to ten years of experience with HPC. Some have Ph.D.'s in science and engineering. These experts bring that experience to customers, if they want it, to make sure their jobs run as fast as possible on POD.

Intel® OPA Enables Non-Blocking Performance on POD

The first POD cluster was built on Intel® Xeon® Processors (codenamed Sandy Bridge and Ivy Bridge) with the Intel® True Scale Fabric host interfaces, supporting Performance Scaled Messaging (PSM), and InfiniBand* (IB) Architecture QDR network (Intel True Scale was the predecessor to Intel® Omni-Path Architecture (Intel® OPA) host interface adapter). From the beginning, all POD clusters run Intel Xeon Processors. The latest deployment utilizes Intel Xeon Processor E5 v4 (codenamed Broadwell) CPUs and Intel OPA.

"We've gained a lot of experience with interconnects over the years," said Pokorny. "We were very excited about Intel OPA in the period before it launched. We've run numerous benchmarks with it, and are pleased with both its performance and scalability within our deployments—both our customer solutions and POD," he concluded.

POD is built on Penguin's Tundra ES scalable cluster configuration. CTS-1 was also built on Tundra ES. The Tundra design is based on the Open Compute Project (OCP) and Open Rack specifications, which use 12-volt power throughout the rack, instead of distributed AC power. "We worked with Intel to design a 12-volt leaf switch for Intel OPA while designing Tundra," commented Pokorny. "We were very glad to get the 12-volt switch. It makes it easy to deploy consistent clusters that only require a power shelf change for a variety of international power requirements."

"But, when customers come to us, they don't ask about the specific components," added Gregorio. "They want to know if we have a non-blocking network, if our systems are fast enough to run their jobs in a reasonable time, and if we have enough cores to support their workloads." The answer to all these questions has been yes.

Results

"Our customers love the POD service," stated Gregorio. "They continue to come back with their jobs, and the customer base is expanding. They like it because it runs on bare metal with consistent and predictable performance and because it's a service. They can budget the core hours they need for the year, and they know what it'll cost them to get their work done."

Penguin recently repeated a series tests that appeared as part of a comparative benchmark of HPC cloud providers that didn't include them. It was posted on arXiv.org in March, and HPCWire covered it. HPCWire then followed up with an article titled, Penguin Takes a Run at the Big Cloud Providers. Their results show they perform better at a lower price than the cloud leaders.

Recently, Penguin has added a small POD cluster built on Intel® Xeon Phi™ Processor with Intel OPA as part of an Intel 'try-before-buy' program. Customers can sign up at the POD service and run jobs on the Intel Xeon Phi Processor- and Intel OPA-based system to evaluate how it will run their workloads. Penguin experts are on hand to help them optimize for the Many Integrated Core architecture of the cluster's processors in order to maximize the performance they can achieve on the Intel technologies.

Solution Summary

Early cloud-based HPC offerings lacked scalable performance. They were expensive to set up and maintain, and they came without expert support. Penguin Computing launched POD as a bare-metal HPC cluster, built on Intel Xeon Processors and Intel OPA fabric, with available software licenses, and predictable billing that was fully supported by HPC experts. Over the eight years POD has been operational, the service has expanded, providing HPC computing performance to a range of industries.

Where to Get More Information

Learn more about Penguin Computing on Demand.

Learn more about Intel® Omni-Path Architecture.

1017/RJMJ/RJ/PDF ♻Please Recycle 336528-001US