END-TO-END SUPPORT TOWARD ACCELERATING DISCOVERY AND INNOVATION WITH A CONVERGED HPC AND AI SUPERCOMPUTER

Produced by Tabor Custom Publishing
in conjunction with:

HPC wire

QCT | intel

Taiwan has been a contributor to major scientific research and creates supercomputers in the region to accelerate discovery and innovation for a variety of industries. At the forefront, is the National Center for High-Performance Computing (NCHC), a national-level research laboratory and Taiwan's primary facility for high performance computing (HPC). NCHC resources include large-scale computational science and engineering, cluster and grid computing, middleware development, visualization and virtual reality, data storage, networking, and HPC-related training.

In order to accelerate the national development of Artificial Intelligence (AI), the Ministry of Science and Technology in Taiwan planned to establish a Taiwan national AI research center. The State Grid Center undertook the construction of cloud services and a big data computing platform for building a Taiwan national AI research and development infrastructure. The goal was to provide cloud integration services for high-speed computing, high-quality networking, high-performance storage, big data analysis and scientific engineering simulation. The goal of the center is to help accelerate AI intelligence in various fields, promote the transformation and leap in related Research and Development technologies and accelerate Taiwan's vision of becoming a "digital country and smart island."

To meet their goals, NCHC needed to develop a supercomputer with compute resources to meet Taiwan's HPC and AI needs. NCHC selected Quanta Cloud Technology (QCT), a global data center solution provider and a long-time Intel® partner, to build their Taiwania 2 and Taiwania 3 supercomputers. QCT is a leading Taiwanese data center company providing end-to-end solutions worldwide that meet their customer's system design needs. Currently, the Taiwania 2 and Taiwania 3 supercomputers are used by academia and industrial users in Taiwan.

NCHC Project Overview

QCT took part in a large scale HPC and AI project, a three-phase project to build powerful, energy efficient supercomputers, and the first AI supercomputer to be entirely made in Taiwan. QCT and partners established the HPC and AI supercomputer implementing a cloud infrastructure for NCHC's projects. QCT's goal was to create a high-density CPU and high-performance GPU computing system to perform AI and HPC operations in virtual machines (VMs) and containerized environments. QCT built a powerful Taiwania 2 AI cluster for NCHC in 2018 as part of their second phase. Taiwania 3 was completed as part of phase 3 in 2020. Both systems include high performance computing capabilities, a large amount of storage space and a secure network built through an AI cloud computing platform. A major focus was also put on energy savings and high energy efficiency.

Taiwania 2 Supercomputer Paves the Future for AI Research in Taiwan

QCT worked with partners to complete the construction of Taiwania 2 to meet HPC and AI supercomputer demands. First of all, AI research requires extensive computational power and Taiwania 2 had to meet this need for Taiwan's researchers. Taiwania 2 was ultimately built with 252 compute nodes, each containing dual Intel® Xeon® Scalable processors and eight state-of-the-art GPUs. The system achieved a Linpack (Rmax) performance of 9 PFLOPS which was faster than the required 7 PFLOPs; it ranked 20th on the TOP500 list of the world's most powerful supercomputers in Nov. 2018, achieving a PUE of 1.16, and ranking No. 10 on the Green500 list of the world's most efficient systems in Nov. 2018. Taiwania 2 is shown in *Figure 1.*

Figure 1. Taiwania 2 supercomputer. Source from NCHC.

QCT also included its HPC and AI Software Stack for rapid deployment, system monitoring and management. Furthermore, its built-in development environment with AI analytics framework and visualization tools include:

- Single-pane-of-glass web-based dashboard to provide real-time monitoring of the system, service status, performance, hardware utilization and health status
- Rapid deployment with automatic software provisioning and configuration of HPC and AI environments

- Dynamic allocation of compute resources to allow users to queue jobs on bare-metal and containerized environments
- Simplified HPC/AI development and simulation workflow with pre-installed and configured tools that include a data visualization viewer, compiler, MPI and AI application frameworks
- Ability to run HPC/AI/data analytics workloads via the web UI or CLI and the remote graphics desktop service

Taiwania 3 Supercomputer

QCT designed Taiwania 3 for NCHC's phase 3 with 900 nodes including a high-density 2U4N HPC CPU cluster, architecture, and support. Taiwania 3 provides ultra-scalability with a HPC CPU cluster using QCT's QuantaGrid T42D-2U, which is equipped with dual Intel® Xeon® Scalable processors and a high-speed network. *Figure 2* shows the Taiwania 3 clusters.

Figure 2. Taiwania 3 supercomputer. Source from NCHC.

QCT performed system integration of Taiwania 2 and Taiwania 3 to ensure compatibility between the systems. Taiwania 2 and Taiwania 3 include high performance interconnects and networking fabric nodes. QCT used a non-Blocking Spine/Leaf architecture to connect all the computing and utility hosts to the InfiniBand interconnect and GPFS parallel file system. Both systems use cabinet water cooling for power efficiency and energy savings.

QCT also provides HPC benchmark capabilities and expertise on industrial HPC/AI benchmarks to ensure that their delivered systems meet performance requirements such as faster simulation times. This allows users to focus on their research rather than worrying about configuration settings.

Taiwania 2: Training a Deep Learning Model to Identify Cancer Metastasis

It is important for diagnosis and treatment of colon cancer to determine cancer stage, which involves the identification of lymph node metastasis. Detecting colorectal tumor metastasis requires a microscopy histologic evaluation of surgical specimens by pathologists, which is time-consuming and prone to misdiagnosis. aetherAI, a leading digital pathology AI solution provider, developed an annotation-free approach to train deep neural networks directly using gigapixel whole-slide images. Trained on Taiwania2, the deep learning algorithm performed well in identifying both macrometastasis and micrometastasis with an area under the receiver operating characteristics curve (AUC) of 0.9993 and 0.9956, respectively.

Taiwania 3: Defending Taiwan — Fighting the Pandemic with Technology

Taiwania 3 was launched in 2020 ahead of its scheduled production date so that it could be used for COVID-19 research as part of the "Tech vs. Virus" project. The "Anti-Pandemic Technology 2.0" NARLabs" NCHC project invited academics, researchers, innovators, and industry leaders to submit proposals to answer the call for COVID-19 virus research. Resources available for the research included the Taiwania supercomputers, Taiwan Computing Cloud (TWCC), DAS data analysis service, and the AIoT platform.

QCT Provides End-to-End Support in Building HPC/AI Systems

Organizations are looking for a best practice system which meets specific workload demands. A workload-driven system capable of running HPC/AI workloads is more important than ever. Organizations face many challenges when building a system capable of running HPC and AI workloads. Furthermore, there are also many complexities in system design and integration. Building a workload driven solution requires expertise and domain knowledge that organizational staff may not possess. In order to simplify the complexity that end users may face, QCT supports their customers and partners in different phases when building workload driven HPC/AI solutions to accelerate their time-to-value.

QCT meets customer needs for creating HPC/AI systems with demand analysis evaluation, architecture design, system deployment, system tuning/benchmarking, and pilot run and implementation:

- QCT provides rapid deployment kits to quickly set up its operating system and HPC/AI environments
- Web-based UI allows administrators to monitor and manage the cluster easily
- QCT Orqestra software allows out-of-band system monitoring for up to 5000 nodes with control by BMC
- QCT AI Labs can do proof-of-concept, testing, tuning, and benchmarking for end users or ISV partners
- QCT has HPC/AI expertise for different scales, different architectures, and domain expertise on benchmarking and tuning
- QCT provides end-to-end support for end users and system integrators to ensure customer success

Building an HPC/AI system requires lots of time and resources–A QCT solution can simplify the process and shorten one's journey from months to days.

QCT Provides State-of-the-Art Servers Using the Latest Intel® Technology

QCT provides state-of-the-art servers powered by Intel® that provide built-in AI and crypto acceleration, Intel® Deep Learning Boost, Intel® Advanced Vector Extensions 512, and Intel® Speed Select that are designed to power new business breakthroughs and scientific discoveries, empowering HPC customers by tailoring performance for their specific workload needs. 3rd Gen Intel® Xeon® Scalable processors offer a balanced architecture with advanced security capabilities with Intel® Software Guard Extension (Intel SGX), Intel® Total Memory Encryption (Intel TME), Intel® QuickAssist Technology, Intel® Crypto Acceleration, and Intel® Platform Firmware Resilience (Intel PFR). These processors in QuantaGrid and QuantaPlex servers are also built for mission-critical, real-time analytics, machine learning, artificial intelligence, high performance computing (HPC), and multi-cloud workloads that support the convergence of HPC and AI in the cloud or on-premises with faster time to results.

QCT servers that support the 3rd Gen Intel® Xeon® Scalable processors include:

QuantaGrid D53X-1U:

- 3rd Gen Intel® Xeon® Scalable processor with built-in AI Accelerators
- Intel® Optane® persistent memory 200 series to propel DAOS and enable AI, HPC and virtualization workloads

QuantaGrid D53XQ-2U:

- 2U server that benefits from DAOS fueled by Intel® Optane persistent memory on the 3rd Gen Intel® Xeon® Scalable platform
- Acceleration card expansion slots to power AI, analytics, bioscience, HPC, and even cloud computing workloads

QuantaPlex T43Z-2U:

- 2U4N high density server optimized for extreme compute performance and space efficiency powered by 3rd Gen Intel® Xeon® Scalable processors

Conclusion

Taiwan is a major contributor to scientific research and has always wanted to expand this capability by creating supercomputers built and located in Taiwan. NCHC selected QCT to build the Taiwania 2 and Taiwania 3 supercomputers to aid in research in Taiwan. QCT is a global solution provider, whose focus is not only supercomputer building in Taiwan, but also providing the capability to build HPC/AI supercomputers worldwide with QCT local offices, QCT AI lab and supportive teams in Europe, the US and APAC.

QCT has expertise in end-to-end HPC, AI and Deep Learning (DL) solutions for companies in a wide variety of fields such as academia, research, life science, and high-technology. QCT creates solutions with different scales, architectures, and software technologies. QCT can provide end-to-end support from system design, through integration, benchmarking and installation for end users and system integrators to ensure customer success.

"QCT is strongly committed to providing our customers with the latest data center technologies," said Mike Yang, President of QCT. "Our collaboration with our technology partners and NCHC on the Taiwania 2 and 3 has produced a wealth of HPC and deep learning solutions that bring the power to academia, governments and enterprises of all sizes in Taiwan. Through joint efforts like these, we are also able to give our customers around the world optimized solutions for different applications, workloads, and business objectives that fit their industry needs."

About QCT

Quanta Cloud Technology (QCT), a global data center solution provider, provide the cutting-edge High-performance Computing (HPC) and AI solutions with vertical applications. QCT solutions combine QCT's advanced hardware infrastructures with software technologies to help build scalable, converged HPC and AI environments across many industries while shortening the time to implement projects so researchers can fully concentrate on their research.

- Learn More: https://go.qct.io/qct-pod/qctpod-for-higher-education-and-research
- QCT Website: https://www.qct.io

HPCwire | QCT | intel