MEET COMPUTATIONAL CHEMISTRY WORKLOADS NEEDS WITH QCT POD SOLUTION

Produced by Tabor Custom Publishing
in conjunction with:

HPCwire

QCT | intel

Executive Summary

Computational chemistry describes the computer modeling and simulation to solve chemical problems. Computational chemistry entails chemists using high-performance computing (HPC) to solve problems by creating models and simulations that require massive amounts of data. An example of using this kind of simulation is identifying sites on protein molecules that are most likely to bind a new drug molecule, creating models of synthesis reactions to demonstrate the effects of kinetics and thermodynamics, and exploring the basic physical processes underlying phenomena such as superconductivity, energy storage, corrosion, or phase changes.

In order to achieve such a goal of chemical activities, chemists would choose the software applications among options. Most applications for computational chemistry are based on open-source such as Quantum ESPRESSO, GROMACS, LAMMPS, GAMESS, NWChem, etc., or license-based software such as Gaussian, Q-Chem, Molpro, VASP, etc.

Moreover, chemists would need to adopt different methods based on Quantum Mechanics during the research. There are several methods that can be implemented in different situations, including using software for Ab-initio, Hybrid-DFT, Semi-empirical and Molecular Dynamics.

Computational Computing Use Cases and Challenges

Computational chemists exclusively work on developing and applying software, collaborating with their colleagues in the laboratory, clinic, or out in the field to apply and validate their models. They may also work with computer scientists who develop advanced hardware and software capabilities for working on especially large or complex problems.

Typical work duties of a computational chemist include:

- Applying software and hardware capabilities for data collection and analysis
- Developing computer models and simulations of chemical and biochemical processes and entities
- Performing and interpreting statistical analysis of large datasets
- Creating known or unknown reaction paths, molecular interactions, or other phenomena
- Designing experiments

There are specific practices within the field that have special HPC requirements. For example, for chemists focused on the creation of new materials, HPC can add another dimension of development of new materials. For instance, the simulation allows materials to be synthesized and characterized virtually. Chemists could use simulation to predict the outcome; such technologies are widely used in Materials Science.

Those seeking to better understand chemical reactions work with different combinations of approaches and materials to find possible reaction paths and they must perform analysis of reaction mechanisms. HPC systems with parallel software could be taken advantage of to discover results.

Those working on experimental drug discovery have yet other requirements. Since there are a huge number of possible materials to synthesize for a target drug, computational aided design has become crucial for drug discovery and design. This is carried out using a combination of Quantum mechanics, molecular mechanics, molecular dynamics on HPC systems.

Regardless the specialization, computational chemists must understand the underlying principles of a simulation, optimization, or other calculation to set up the conditions and parameters of their study and to ensure that the results are meaningful and properly interpreted. And they run progressively more sophisticated models, simulations, and analytic routines. These efforts require vastly more compute power than has ever been needed in the field before. Additionally, speed to results and insights is critical to all areas and applications when conducting such research.

One other factor is the complexity of these systems. In the past, the work was the domain of larger institutions and commercial enterprises that had experts in software development, hardware maintenance, system administration, and modeling applications. Now, smaller companies and academic departments require the same HPC capabilities, delivered as easy-to-use complete solutions that require little or no ongoing system management on the part of the researchers.

QCT Solution for Computational Chemistry

QCT has developed a pre-configured and pre-validated platform powered Intel® Xeon® Scalable Processors to speed time to discovery for scientific research, namely QCT POD for Higher Education and Research. It provides best practice hardware and software integration for computational chemistry to enable parallel computing capabilities. QCT HPC experts help design best-fit compute, storage, and networking building blocks for their computational chemistry customers.

QCT POD, which stands for platform on-demand, provides the infrastructure and software stack needed to meet the processing and storage requirements of modern research. Furthermore, the QCT POD for Higher Education and Research offers an innovative technology system with building blocks designed to meet different research demands, including those found in quantum chemistry and computational physics.

It offers an optimized run-time environment for HPC/AI workload applications for a streamlined HPC/AI development environment and workflow. Additionally, the system includes cluster management software, built-in application frameworks, compute-optimized hardware, storage technologies, and low-latency networking, all seamlessly integrated and validated to accelerate the HPC research process.

Compared to assembling a home-grown solution, using the QCT POD simplifies the journey from system preparation to deployment phase and system management and monitoring. Once it is installed, the support continues with QCT HPC experts who are available to offer system configuration consultation and benchmarking and tuning service on a QCT Solution Center to ensure the quality and serviceability of the infrastructure.

Using the QCT POD provides benefits to different stakeholders. For example:

- End users achieve faster productivity and time to results, getting deeper insights faster.
- System administrators get a solution with built-in system management and monitoring, removing the burdens of the many chores they would otherwise need to conduct to keep the system running.
- Developers receive a pre-built and comprehensive environment to do their work, eliminating the need to assemble a system every time a new application is created.

Taking a Deeper Look at the QCT POD for Higher Education and Research

Diagram 1 demonstrates the building blocks of QCT POD. The system consists of compute, storage, management, and networking blocks. It also includes the QCT HPC Workload Package to speed a researcher's time to simulation and the QCT POD Manager to provide a user-friendly dashboard and ease of monitoring and management.

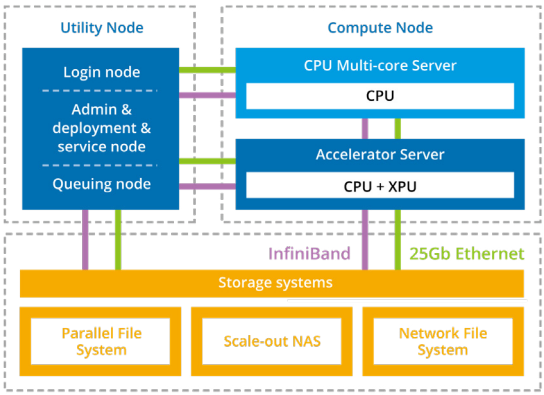

Diagram 1

Taking a deep look at the QCT POD suite in diagram 2, QCT defines four levels to meet researchers' demands and workloads. QCT POD includes QCT's best-in-class hardware and an integrated software and system environment to allow researchers to achieve "simplicity" and "efficiency" all the way from system planning and preparation, implementation, operation, and maintenance.

Building Blocks			Starter		Standard		Advance		Premium	
			Qty	Specification	Qty	Specification	Qty	Specification	Qty	Specification
Utility Node (Login, admin, service and queuing)	All in one utility& storage node		1	QuantaGrid D53XQ-2U	1	QuantaGrid D53XQ-2U				
	All in one Utility Node						1	QuantaGrid D53X-1U	1	QuantaGrid D53X-1U
Computing Node	CPU Node		4-6	QuantaGrid D53X-1U	6-8	QuantaGrid D53X-1U	6-8	QuantaGrid D53X-1U	6-10	QuantaGrid D53X-1U
Storage System BeeGFS Filesystem	Metadata Server						1	QuantaGrid D53XQ-2U	1	QuantaGrid D53XQ-2U
	Object Storage Server						2	QuantaGrid D53XQ-2U	2	QuantaGrid D53XQ-2U
	Storage Enclosure for Object Storage Target						1	QuantaVault JB4602	1	QuantaVault JB4602
Networking	Data& Management Network	25G Ethernet Switch	1	QuantaMesh T4048-IX8D	1	QuantaMesh T4048-IX8D	1	QuantaMesh T4048-IX8D	1	QuantaMesh T4048-IX8D
	Management	Management Switch			1	QuantaMesh T1048-LB9M			1	QuantaMesh T1048-LB9M
	Data Network	Infiniband Switch			1	HDR Infiniband Switch			1	HDR Infiniband Switch

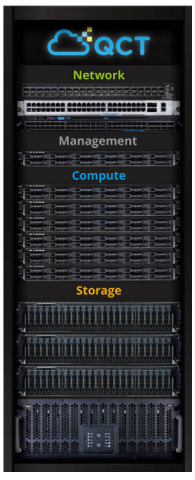

Diagram 2

With regard to compute and storage capabilities, QCT adopts in-house hardware technologies as below:

QuantaGrid D53X-1U:

A Balanced architecture 1U compute server each equipped with two 3rd Generation Intel® Xeon® Scalable Processors, which optimize for AI, HPC, data analytics workloads, supporting 2nd Gen Intel® Optane™ PMem 200 Series Barlow Pass.

QuantaGrid D53XQ-2U:

Designed with scalable and flexible NVMe, each equipped with two 3rd Generation Intel® Xeon® Scalable Processors, to support for HPC/AI storage nodes and deliver high performance for a parallel file system or scale-out NAS. QuantaGrid D53XQ also can support 2nd Gen Intel® Optane™ PMem 200 Series Barlow Pass to boost memory capacity.

Storage System

QCT offers different storage systems based on workloads and user demands, as shown in diagram 1. In terms of the file system, QCT offers a BeeGFS Parallel file system to enable easy-to-use features and is a great fit for I/O intensive HPC workloads. Through such experiences, QCT offers a system that is:

- **Easy to deploy:** the setup process of BeeGFS has been streamlined and automated through QCT Rapid Deployment Kits.
- **Optimal architecture design:** the most optimal BeeGFS design in terms of hardware design, software architecture, and configurations to deliver optimal performance.
- **Unified management:** a GUI dashboard to perform administrative management tasks and to monitor the state of the file system and its components.

Deployment and Management Made Easy

The QCT POD for Higher Education and Research is simple to deploy. The solution features automatic offline system deployment and scaling of bare-metal environment, automatic software provisioning and configuration, and a built-in software repository for a bare-metal and containerized environment. Managing these aspects of the system frees up IT staff that would normally perform these tasks.

Once a system is installed, additional built-in features ensure that it keeps running with minimal IT intervention. QCT POD Manager offers real-time system monitoring and simplified cluster system management to streamline administration workflow with a QCT web-based GUI dashboard.

Ready to Use HPC Applications

The QCT POD includes the QCT HPC Workload Package, which is a package of pre-compiled and pre-configured HPC applications that are delivered in the form of environment modules to allow researchers and developers to run simulations instantly.

The QCT HPC Workload Package provides various types of HPC applications that are commonly used in different research fields.

Using its HPC expertise and computational chemistry domain knowledge, QCT pre-compiles and pre-configures the HPC applications with optimized and best-practice configurations. Managing these tasks means researchers and developers don't have to compile and configure these applications from ground zero and go through a trial-and-error process to find optimized configurations.

QCT also manages software dependencies. In a typical computational chemistry HPC site, there are multiple researchers and users working on different research projects. Each of these projects runs different applications that are dependent on different compilers, libraries, and MPIs to work. Researchers and developers often spend quite a bit of time to ensure the software dependencies are set properly before actually running simulations. For admins who have to get these applications ready for users, installing the application and managing the dependencies is a complex and time-consuming process.

With QCT's Engineering Team's experiences and domain knowledge in these applications, QCT has done extensive work to manage the complex software dependencies to build QCT HPC Workload Package for users and developers, allowing them to use the software without worrying about the dependency issues.

Real-World Proof Point

QCT accumulates numerous experiences in building Supercomputer Taiwania2 and 3 in HPC/AI research field to ensure the reliability, serviceability, and optimized performance of the solution to assist a national laboratory and academic institutions in solving complex problems in different research fields such as weather forecasting, imitating atom's movement and interaction or understanding the electronic structure, astronomy research, new material development, and drug discovery.

QCT built the Taiwania 2 supercomputer for a national laboratory. QCT established an AI cloud computing platform, which could provide real-time and convenient computing services to the industry, universities, and institute circles. QCT delivered an AI cluster that includes x252 compute servers ranked on the Top 500 list of the world's most powerful supercomputers, each equipped with dual CPUs and 8 GPUs. Additionally, the system has direct-to-chip liquid cooling to enable superior energy efficiency to place it on the Green 500 list.

QCT built another Taiwania 3 supercomputer, adopting ultra-high-density 2U4N compute nodes, with 900 compute nodes in the cluster that also ranked on the Top 500 list. For the large number of compute nodes, QCT introduced QCT Rapid Deployment Kits to simplify the system deployment journey and provide custom built-in software environments to accelerate time to market.

Working with a Technology Partner

QCT develops solution centers at headquarters across hundreds of nodes to streamline systems through configuration, benchmarking, and the certification process to industrial standards to ensure key workload performance.

In addition, QCT has HPC/AI experts in molecular dynamics, chemistry, weather and climate, artificial intelligence, and next-generation sequencing, allowing it to provide professional benchmarking and tuning service upon customer request.

To learn more about QCT POD for Higher Education and Research, visit: https://go.qct.io/qct-pod/qctpod-for-higher-education-and-research/

About QCT

Quanta Cloud Technology (QCT) is a global data center solution provider. We combine the efficiency of hyperscale hardware with infrastructure software from a diversity of industry leaders to solve next-generation datacenter design and operation challenges. QCT serves cloud service providers, telecoms, and enterprises running public, hybrid, and private clouds.

Product lines include hyper-converged and software-defined data center solutions as well as servers, storage, switches, integrated racks with a diverse ecosystem of hardware components and software partners. QCT designs, manufactures, integrates, and services cutting-edge offerings via its own global network. The parent of QCT is Quanta Computer, Inc., a Fortune Global 500 corporation.

http://www.QCT.io

—