QCT UNIQPOD FOR HEALTHCARE:
AN OPTIMIZED CONVERGED AI/HPC PLATFORM TO REALIZE AI MEDICAL IMAGING

Produced by Tabor Custom Publishing
in conjunction with:

HPC wire

QCT | Qumulo | NVIDIA

Executive Summary

There is massive growth in the number of medical images which must be processed, analyzed, and stored. This growth is placing a burden on medical staff such as Radiologists who must analyze the images. Artificial Intelligence (AI) is increasingly being used to analyze medical images generated from a wide variety of medical devices. There are many challenges in adopting an appropriate medical imaging strategy to meet today's medical processing, storage, and AI training and analysis needs.

This paper provides information on how the NVIDIA Clara AI solution aids in developing, training, and managing AI medical imaging applications. In addition, it describes how QCT's UniQPOD for Healthcare solution along with the integrated Qumulo storage provides the compute and storage infrastructure needed in various AI applications in the medical field. The UniQPOD for Healthcare solution automates launching NVIDIA Clara and allows users to easily manage all NVIDIA Clara workflow tasks within the UniQPOD for Healthcare workspace and provides powerful and easily-managed infrastructure to empower NVIDIA Clara.

Challenges of Analyzing Medical Images in Radiology

The growth of medical images in the field of Radiology provides a good example of the challenges faced in analyzing medical images. The advances in Magnetic Resonance Imaging (MRI), Magnetic Resonance Enterography (MRE), X-Rays, and Computed tomography (CT) scans produce thousands of high-definition images that must be analyzed by Radiologist. These images present challenges for Radiology departments resulting in too much work for existing Radiology staff. In addition, there may be a shortage of Radiologists trained in analyzing the specialized images. With the deluge of images to evaluate, there is also a higher possibility of a delay in image evaluation or making errors in analysis leading to poor medical outcomes.

A promising area in healthcare innovations is using Artificial Intelligence (AI) to analyze medical images such as those typically analyzed by Radiologists or to discover anomalies not discernable by the human eye. However, building an effective AI model with high accuracy is time consuming as well as data intensive. A Radiologist or AI researchers must annotate images to denote information such as abnormalities managing AI workloads during production requires extensive work because there are many different device data sources which increases the complexity to integrate these siloed devices to build AI pipelines.

Challenges of Storage Systems for Implementing AI in Medical Imaging

Medical devices like MRI, CT scans, and IoT devices are producing medical images with larger file sizes and numbers as well as a variety of file types. Effective analysis of medical image data first requires a comprehensive and sophisticated analysis and then this data must be effectively stored. AI training requires very fast access to the data with high bandwidth and low latency. Traditional storage systems fail to meet the performance required for AI training and scalability to handle and store the large amount of data, unstructured data, various file sizes, and files protocols. Files of different sizes present unique challenges for the storage systems to provide the high performance required and store the data efficiently for cost reasons.

NVIDIA Addresses Challenges in Applying AI to Medical Imaging

Though the power of AI can bring innovation to medical imaging, medical organization still face great challenges in training AI models and managing the complex workflow. NVIDIA Clara is purposely-built to address these challenges. The NVIDIA Clara tools have greatly advanced the development of medical imaging AI applications framework using PyTorch and MONAI frameworks which are meant to achieve zero coding for DL models and training. Clara imaging and NVIDIA NGC offers tools like AutoML, transfer learning, AI-assisted annotation for labeling datasets, federated learning and over 20 pre-trained AI models to get started quickly. Clara includes full-stack GPU-accelerated libraries, SDKs, and reference applications for developers, data scientists, and researchers to create real-time, secure, and scalable solutions.

These Clara Train Application Framework tools aid developers and researchers in speeding up the training of AI models:

- AI-Assisted Annotation applications: This tool contains different types of pre-trained AI models to help Radiologists automate and speed up the image annotation AI process.
- Pre-trained models specific for medical imaging: Users can pull different pre-trained models to leverage transfer learning techniques to greatly accelerate the time-consuming AI training process.
- AutoML: AI-assisted AI model tuning can suggest parameter tuning configurations for AI models.

The Clara Deploy SDK includes features that facilitate the process of deploying a trained AI model. Clara Deploy includes an AI pipeline template that allows the developer or operator to pull different AI pipeline templates to facilitate the deployment process. A built-in Digital Imaging and Communications in Medicine (DICOMM) adaptor allows Clara to be integrated with most medical devices, such as equipment used in MRE or CT scans, to access data directly from the source to save time. An integrated Workload manager allows operators to view resource utilization and automate scheduling, queuing, and provides prioritizing of multiple AI inference workloads.

How QCT Supports Medical Image Processing, Storage and AI Use

While NVIDIA Clara plays a critical role in helping medical institutions and organizations to unleash the power of AI in medical imaging, the underlying infrastructure to support this powerful application also has equal importance. UniQPOD for Healthcare is the easily-managed underlying infrastructure that medical institutions and organizations need to provide optimal compute power, storage, and networking to empower NVIDIA Clara. UniQPOD for Healthcare also contains purposely-built software features that automate and streamline the workflow of NVIDIA Clara.

QCT is a leading cloud datacenter solution provider with extensive experience in developing High Performance Computing (HPC) and AI solutions for companies in a variety of fields. UniQPOD for Healthcare is a QCT Platform on Demand (POD) solution, which offers an innovative technology system with building blocks designed to meet different demands of AI application and data analytics in the medical and life science field, like Next-Generation Sequencing (NGS), Molecular Dynamics (MD) and Medical Image Recognition. QCT integrates a Qumolo storage software which allows storage of files typically produced during medical imaging. QCT ensures the quality and serviceability of infrastructure and has innovative features to make it easy for users to launch and use NVIDIA Clara.

UniQPOD for Healthcare Meets Image Processing and AI Needs

Figure 1 shows the architecture of the UniQPOD for Healthcare solution that provides dedicated management services along with a powerful compute infrastructure capable of meeting diverse workload demands and medical image processing needs. UniQPOD for Healthcare storage provides data lifecycle management for storage system scalability and data protection.

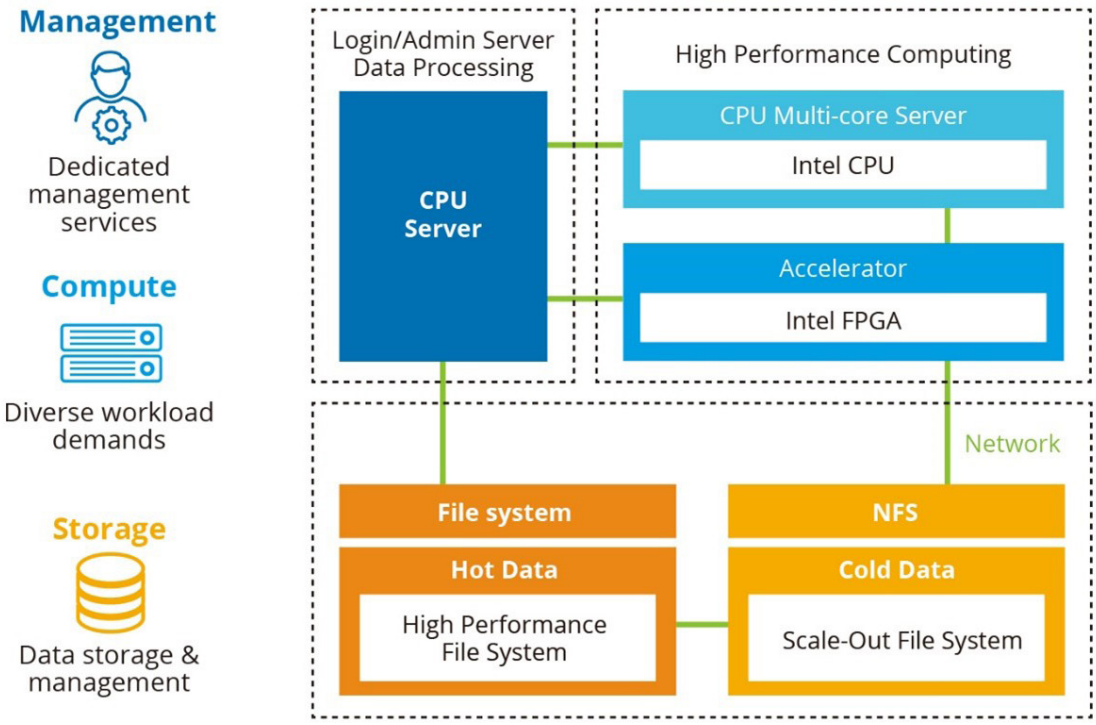

Figure 1. QCT UniQPOD for Healthcare Architecture

UniQPOD for Healthcare System Custom Designed for Different Workloads and Demands

QCT developed the UniQPOD for Healthcare Starter/Standard-Suites and UniQPOD for Healthcare Advanced/Premium-Suites to meet the varied needs of users with different performance, capacity, scalability, and cost concerns. The UniQPOD for Healthcare Starter/Standard-Suites are all-in-one cost effective and compact system with a small footprint which provides all necessary compute, storage, networking, and software frameworks to run AI workloads. It is an ideal solution for small and local medical institutions or organizations to develop and run AI applications. The UniQPOD for Healthcare Advanced/Premium-Suites provide large medical institutions, organizations, or life science research centers with an enhanced performance solution which allows users to fully unleash the power of AI and HPC in diverse use case scenarios. The Enterprise Suite provides a performance-oriented and capacity-oriented option for shared storage to provide flexibility to meet different storage requirements. Figure 2 shows the QCT hardware and building blocks in the Starter/Standard and Advanced/Premium suites.

Building Blocks		Starter		Standard		Advanced		Premium	
		Qty	Specification	Qty	Specification	Qty	Specification	Qty	Specification
Management Nodes	Deployment & Login	1	D43N-3U	1	D43N-3U	1	D43K-1U	1	D43K-1U
	Admin & Service					1	S43CA-2U	1	S43CA-2U
Computing Nodes	HPC CPU					1	D43N-3U	1	D43K-1U
	AI training							2	D43N-3U
Storage Nodes	Capacity			1	T22P-4U2N	1	T22P-4U2N*	2	T22P-4U2N*
	Performance					2	D53XQ-2U*	4	D53XQ-2U *
Network Fabrics	Out-of-Band Management			1	T1048-LB9M	1	T1048-LB9M	1	T1048-LB9M
	Service Management					1	T4048-IX8D	1	T4048-IX8D
	Compute/Storage Network								

Notes:
AMD server marked in Orange
Intel server marked in Blue
* Please consider either capacity or performance maximized storage nodes for Enterprise versions

Figure 2. UniQPOD for Healthcare Laboratory and Enterprise Suites

Benefits of UniQPOD for Healthcare Software Suite

The built in UniQPOD for Healthcare Software Suite offers a full-suite of services that provides System Administrators with effective system monitoring and management tools as well as benchmarking, tuning and diagnostic tools. Developers and end users have access to development tools including compilers, libraries, application frameworks and data visualization tools. Additionally, UniQPOD for Healthcare is optimized to handle HPC and AI workloads.

UniQPOD for Healthcare Automates and Streamlines the NVIDIA Clara Workflow for Users

Launching NIVIDIA Clara requires manual preparation of containers, starting servers, installing image viewers, and sending data between servers. QCT has done extensive work behind the scenes to streamline and automate the setup process to save time and effort for researchers and developers to allow them to just focus on leveraging NVIDIA Clara to build their AI applications.

QCT Automates NVIDIA Clara Process with a Few Clicks

Using UniQPOD for Healthcare, with only a few clicks on a Jupyter Hub link, the required NVIDIA AIAA server and MITK client are prepared and ready to use. The user does not need to worry about the data transfer between the client and server since all the data is stored in UniQPOD for Healthcare storage and is accessible by both. This solution makes it easier for both users and developers and solves many AI workload and data challenges. In addition, the whole process—allocating resources, preparing the data model, data preprocessing, and training—can be performed in a single workspace on the UniQPOD for Healthcare solution. QCT also provides unified management of compute and storage from a single pane of glass dashboard. QCT understands the workload behavior of medical imaging and offers best-practices and pre-defined hardware configurations with UniQPOD for Healthcare installed.

Figure 3 shows an example of the UniQPOD for Healthcare automated NVIDIA Clara process.

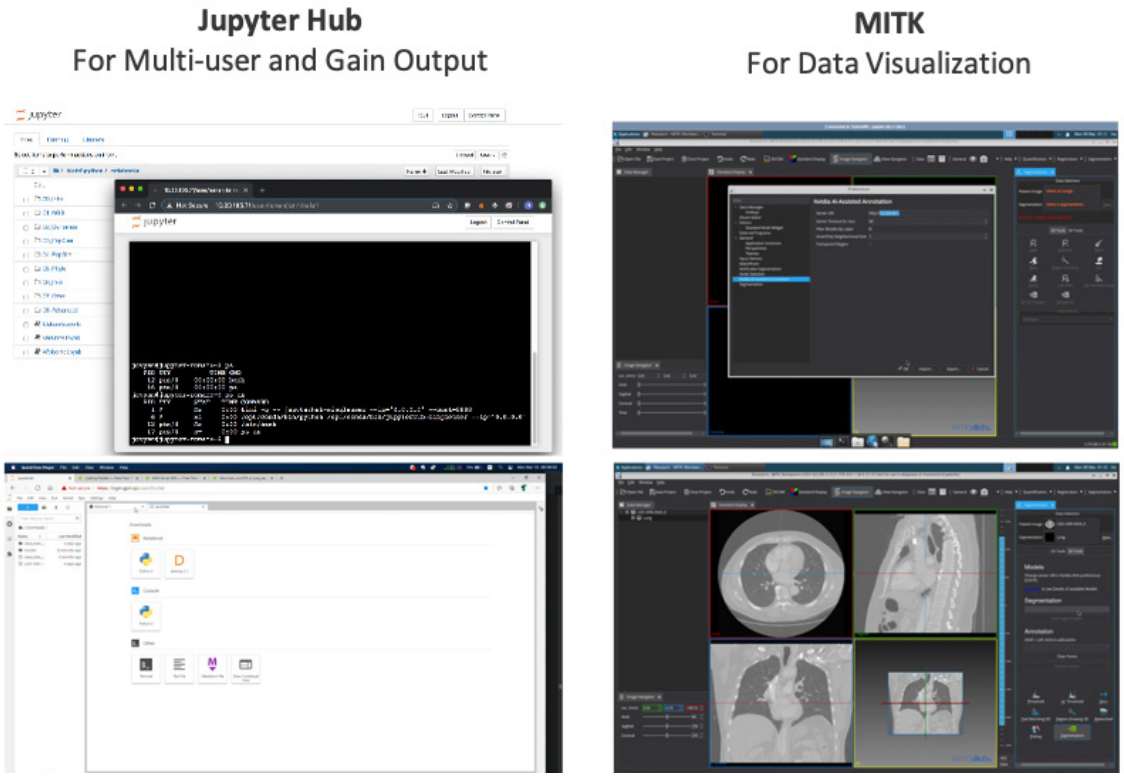

Figure 3. UniQPOD for Healthcare automated NVIDIA Clara medical imaging process.

How Qumulo Addresses Medical Image Storage Challenges

Qumulo is a high performance storage system that provides the high bandwidth and low latency required for AI training. Qumulo's file data storage solution can integrate perfectly with QCT solutions to allow organizations to easily store and manage medical imaging data. The Qumulo high-performance data storage solution works with unstructured data, many different file types produced during the medical image processing and provides easy management, scalability, and data tiering.

Qumulo provides the high bandwidth and low latency required for AI training using predictive caching and proactive prefetch of data. Implementation of common segment allocation provides optimized access to small and large files. Qumulo software contains real-time analytics to identify hotspots and display throughput trending.

As MRI, CT scans, X-rays are producing medical images of higher resolution with larger file sizes, storing massive amount of data effectively and maintaining scalability for exponential data growth in the future becomes a problem. Qumulo has the capacity to scale and manage a large number of files and data in a cost-effective manner. Qumulo's architectural design supports hundreds of nodes, up to 18 quintillion files, file sizes of up to 9 exabytes, and there can be up to 4.25 billion files in an individual directory. Adding more nodes, each with its own storage capacity, is a seamless operation with data automatically distributed. Data is moved to more economical hard disk drives when it isn't being actively accessed—all within a single tier.

Summary

Medical tests such as MRE, MRI and CT scans produce massive amounts of high-resolution images that must be analyzed for a medical diagnosis. Radiologists typically analyze and make a diagnosis from these images, but image growth is placing a burden on Radiology staff to provide a timely and accurate diagnosis. Artificial Intelligence (AI) is increasingly being used to analyze medical images to provide a machine-assisted diagnosis. Developing and maintaining AI applications for medical image analysis and diagnosis is complex and presents many challenges.

The NVIDIA CLARA tools make it easier for organizations to handle the entire AI workflow.

QCT developed the UniQPOD for Healthcare solution containing medical workload building blocks, a streamlined deployment process, and a simplified system and storage management solution for medical clients. UniQPOD for Healthcare provides researchers, developers, and medical organizations with a streamlined way to launch and use NVIDIA Clara tools. UniQPOD for Healthcare allows users to easily manage all NVIDIA Clara workflow tasks within the UniQPOD for Healthcare workspace.

For more information on QCT and how UniQPOD for Healthcare can help your organization, see:

QCT Platform on Demand: https://go.qct.io/qct-pod/

UniQPOD for Healthcare: UniQPOD for Healthcare | QCT