FRONTIER ENTERPRISE

WHITEPAPER

Building the Essential Foundations of Smarter Healthcare

Published by Jicara Media in partnership with Lenovo and Intel.

A common refrain in the business media is that the global COVID-19 pandemic has highlighted the importance and urgency of digital transformation across all organizations. This is true, but even before the pandemic emerged, the healthcare sector was arguably already more aware of this than most.

According to a recent report from McKinley & Co, the Asia-Pacific healthcare industry has been increasingly under pressure from several fundamental forces driving the need for change, including: a growing elderly population, supply constraints to meet demand, the growing financial burden of healthcare and rising consumer expectations.

Those forces are driving the development of customer-centric healthcare ecosystems revolving around use cases such as broadening access to primary care, expediting access to acute care, managing and monitoring disease and improving health and wellness as a preventative measure to catch potential health problems early before they escalate into something expensive to treat.

THE FUTURE OF SMART HEALTHCARE

The details may depend on which ecosystem partner you talk to, but the overarching vision is one where healthcare providers can monitor our vital statistics and health lifestyles by remote – what we eat, how much exercise and sleep we get, etc – and spot warning signs of possible illness early before symptoms evolve into something serious. We can consult doctors online and receive medication by delivery, making consultations less of a hassle. For people who have already been treated for illness, doctors can monitor their recovery from home. And diagnoses themselves are aided by powerful artificial intelligence algorithms to make the process faster and more accurate.

Much of this is already happening to varying degrees in countries like China, Singapore, Indonesia and India. For example, Singapore's HealthHub application enables patients to book appointments, order medications, and access vaccination records and disease risk assessments, among other things. In China, Medlinker and ClouDr are building platforms to connect patients, physicians, payers, providers, and pharmacies to help manage chronic illnesses such as diabetes and kidney disease.

And this is just the start. According to a recent survey from PwC's Strategy&, healthcare executives firmly believe that by 2035, today's healthcare system will evolve into an intelligent system in which healthcare is personalized, digitized, preventative and seamlessly integrated into daily life.

This will require all stakeholders in the healthcare ecosystem to have a deeper understanding of individual human biology and lifestyle, to be able to deliver disease and well care earlier and in a more personalized manner – and that will require access to a vast array of useful data on things such as fitness, nutrition, and social determinants of health to as multi-omics and disease progression and expression.

Data will also play a core role on the diagnostic side of healthcare – such as early detection of tumors, for example – as well as the treatment side as bio-pharma companies leverage data to develop personalized health solutions such as drugs and digital health applications.

Indeed, for just about every aspect of the future of smarter healthcare, data is the most essential common component.

Consequently, the foundation of smarter healthcare comprises two crucial technologies:

- Artificial intelligence algorithms to make sense of all that data and produce insights
- A high-performance computing (HPC) platform powering that AI.

APPLIED AI – MACHINE LEARNING AND DEEP LEARNING

AI is becoming increasingly ubiquitous in the business world, but the starting point is to understand that AI is not just one technology. There are actually several different flavors of AI, most of which can be used in the healthcare field, but each works in different ways – in other words, the specific processes and tasks they can support vary widely.

For example, machine learning (ML) – one of the most common forms of AI – is a statistical technique that trains specific models with focused data. ML is a broad technique at the heart of many approaches to AI – so much so that even ML comes in different flavors.

In healthcare, the most common application of traditional ML is precision medicine – the prediction of which treatment protocols are likely to succeed based on learned knowledge of prior patients and the treatments used.

Put simply, ML takes a database of learned knowledge and applies ‘rules’ based on things like, say, a type of drug used with a certain type of disease on a particular age group or gender. The ML processes the data within those parameters to predict a potential/probable outcome based on that criteria. This not only frees the doctor from having to scour databases from all over the world across multiple volumes of prescription data and results data, etc, but also enables the results to be obtained much faster. That in turn enables patient care to begin much earlier.

A more complex form of machine learning is deep learning – neural network models with many additional levels of features and variables that can predict outcomes. There may be thousands of hidden features in such models, which only today can be discovered and used because we have the ability to use fast-processing CPUs and GPUs.

A common application of deep learning in healthcare is the recognition of potential cancers from imagery such as MRIs and CT Scans. Currently, vision-based deep learning is typically used in oncology image analysis, and promises greater accuracy in diagnosis than the previous generation of automated tools for image analysis.

As an example, in 2017 Lenovo Research developed a solution called E-Health that uses automatic segmentation algorithms to segment liver lesions in contrast-enhanced abdominal CT scans. The segmentation of CT scans plays a critical role in the diagnosis of diseases and contributes significantly to clinical treatment and pathology research.

E-Health uses deep learning to analyze the CT images of patients to detect tumors. If a tumor is found, the algorithms provide crucial information about them such as size, location and type, which can assist doctors in diagnosis. E-Health also integrates the basic information of CT images and patients and annotates the locations and status of tumors, making it easier for doctors to read CT images. The solution is also able to display the tumor as a 3D model integrated with relevant information about the tumor, which makes it easier for doctors to explain the diagnosis to their patients, bringing more humanity, compassion, and trust to a difficult moment.

Obviously this application of deep learning can extend well beyond image analysis. For example, a number of startups are focusing specifically on diagnosis and treatment recommendations for certain cancers based on their genetic profiles. This is significant because many cancers have a genetic basis – consequently, clinicians have found it increasingly complex to understand all genetic variants of cancer and their response to new drugs and protocols. AI is absolutely essential to wade through the rivers of data being produced by genomics sequencers.

POWERED BY HPC INFRASTRUCTURE

That brings us to high performance computing, which is responsible for crunching all that data quickly. In fact, HPC infrastructures are essential for supporting the compute demands of machine learning and deep learning software suites.

Supercomputers are already a fixture in the healthcare sector for research purposes. Pharmaceutical companies, universities and government research labs, all use HPC to do research, whether that's analyzing a whole human genome sequence or sequencing the SARS-CoV-2 coronavirus and developing a vaccination model.

The traditional image of supercomputers is that they're physically huge and expensive, but this is changing fast. Moore's Law is now doubling compute power every seven months instead of every 12 to 18 months. Meanwhile, the ability to fabricate silicon substrates has increased dramatically over the last decade, enabling chip sizes to shrink down to 7 nm and beyond. All of this is bringing us into the GPU era that provides the required power to compute AI/ML applications.

Moreover, smaller HPC infrastructures means lower costs. For example, today's genome sequencing instruments generate up to 6 Terabytes of data in 24 hours each. Sequencing 60 human genomes in one run generates a massive amount of data that requires faster data analysis and storage to keep up. It's also time-consuming – sequencing one sample from one person takes anywhere from 40 to 150 hours, which is neither scalable nor efficient for a genomics study.

It's also expensive – or at least it used to be. According to the National Human Genome Research Institute, in 2007 it cost around $10 million to sequence just one genome. Thanks to the advances in HPC infrastructure, the cost is now well under $1,000, and is projected to go down even further.

Moreover, the cost of HPC equipment itself is dropping as form factors get smaller – what costs millions of dollars today will be reduced to tens of thousands of dollars for the same amount of compute power. That will also enable supercomputers to move from the heavy-duty research centers to the front offices, where more companies in the healthcare ecosystem will have affordable access to that level of computing power and storage to harness AI for their own use cases and research.

SPEED, EFFICIENCY AND ACCURACY

One caveat with AI is that it's not a magic box – dep learning algorithms have to be designed properly to make predictions or recommend treatment courses, and the data itself has to be clean and reliable. This is absolutely crucial in the healthcare field, where decisions are literally a matter of life and death. That means the speed, efficiency and accuracy of these algorithms – as well as the math that goes into these algorithms and the data scientists that associated with it and their skill sets – becomes incredibly important.

For the speed component, there are a couple of dimensions to this. For example, Lenovo and Intel have developed a technology suite for processing genomics data called GOAST (Genomics Optimization And Scalability Tool), which reduces the sequencing time of a whole human genome from 40-150 hours to less that 1 hour. The other speed component is the ability to share data and results quickly, which GOAST also enables.

That ability to process and share data fast also enables the efficiency part of the equation. Indeed, it's one reason why researchers were able to generate a COVID-19 vaccine quicker than people expected. The SARS-CoV2 coronavirus is similar in genetic makeup to the SARS virus, yet fundamentally different. So the ability to sequence the coronavirus quickly was incredibly important. It also required a high level of collaboration with multiple organizations sharing that data around the world, highlighting the importance and value of enabling data to be shared quickly.

Finally, the accuracy component relies on collaboration with medical experts to develop the algorithms and ensure that the results they produce check out.

CASE STUDY

Powering groundbreaking plant genomics research.

How the University of Delhi Department Genetics uses a Lenovo high-perfomance computing (HPC) system, featuring 2nd Gen Intel® Xeon® Scalable processors, to reduce time to insight.

Lenovo Infrastructure Solutions for The Data-Centered

1

Background

The University of Delhi is one of India's largest and most renowned higher education institutions, with a community of over 62,000 students. Since it was established in 1984, the university's Department of Genetics has emerged as a major hub for teaching and research in genetics, genomics, and molecular biology.

The department is recognized internationally for its scientists' work on plant and human genetics, plant breeding, crop improvement, and crop health. The human geneticists at the University of Delhi perform Omics (e.g. genomics, transcriptomics) analytics to understand the genetic basis of several human diseases, including schizophrenia, Parkinson's, and rheumatoid arthritis. The Centre for Genetic Manipulation of Crop Plants (CGMCP) was established in 1996 at the South Campus of the University of Delhi with funding from the National Dairy Development Board (NDDB).

Since its inception, a major mandate of CGMCP has been to improve the productivity of oilseed mustard – *Brassica juncea*. This crop is grown extensively in the dryland areas of the Indian subcontinent as a source of edible oil for human consumption and seed meal for livestock feeding.

Recently, scientists at CGMCP started using Omics approaches to understand plant-pathogen interactions and to understand factors underlying in high yield of the crop for different varieties of the species.

2

Challenge

With dozens of pioneering new research projects starting up at the University of Delhi's Department of Genetics each year, demand for high-performance computing (HPC) resources is extremely high. The department operates a central HPC cluster that all members of the faculty can access—but there was often a long wait for essential compute resources.

To accelerate vital research on oilseed brassicas, CGMCP decided to invest in its own dedicated HPC infrastructure. CGMCP sought an affordable solution that would deliver the high performance and high throughput capacity required to support large-scale genomics analytics.

Cutting-edge HPC.

CGMPC engaged Lenovo to design and deploy an HPC solution that met its demanding performance and memory requirements.

Lenovo proposed the Genomics Optimization and Scalability Tool (GOAST) Plus, an HPC architecture specifically optimized for genomics analytics. The pre-configured hardware and software bundle, based on Lenovo ThinkSystem SR950 servers with 2nd Gen Intel® Xeon® Scalable processors, is extremely fast, affordable, and easy to use.

Dr. Paritosh recalls: “We were very impressed by the technical specifications of the proposed solution. The ThinkSystem SR950 is an eight-socket server, and memory capacity is an important consideration for genomics analyses. We are impressed by the performance of the system for de novo genome assembly with data generated from the third-generation sequencing approaches.”

Why Lenovo? All-in-one HPC for genomics.

Working closely with Lenovo, CGMPC installed two Lenovo ThinkSystem SR950 servers, each equipped with eight Intel® Xeon® Platinum 8260 Processors, and a Network File System (NFS) based on a Lenovo ThinkSystem DE2000H Hybrid Storage Array.

The GOAST system includes the CentOS operating system, the Genome Analysis Toolkit (four different types of variant-calling workflows in the GATK) and the software suite from the Broad Institute, as well as all other software dependencies required to run GATK.

The Lenovo team configured and optimized the GOAST system to deliver optimized performance, tuning the hardware so it can run at full capacity when performing genomics analytics. Lenovo also created easy-to-run scripts that wrap around the complex and often

difficult-to-set-up GATK workflows to simplify submission, monitoring, and management of genome samples thus greatly improving usability for researchers.

3

Result

With the Lenovo GOAST system in place, scientists at CGMCP have access to powerful HPC resources that are extremely fast, affordable, and easy to use. Crucially, high throughput (more samples analyzed at any one time) and short execution times are accelerating time to insight.

One key area of research at CGMCP is a genetically modified variety of mustard that has demonstrated increased yields over existing varieties. To reduce India's need for edible oil imports, scientists at CGMCP are working to further improve this mustard variety, using insights from the Lenovo HPC solution. For example, while identifying and manipulating genes related to drought and disease resistance to breed hardier crops, Dr. Kumar could also map, tag, and introgress traits in *Brassica juncea* for resistance to *Albugo* candida – a type of white rust fungus that can cause serious damage to the plants.

Dr. Paritosh confirms: “I can now process more genomes concurrently and get results faster. Previously, it took 48 hours to process a ~110X *Brassica juncea* genome. Now it takes just six hours, thanks to GOAST’s 8x boost in performance for ultra-deep genomes.”[1]

✓ **36x and 8x increase in performance for human and plant genomics analytics, respectively[2]**

✓ **1.3 hr. to process a 30X human whole genome; 6 hours to process a 110X *Brassica juncea* whole genome[3]**

✓ **Faster time to insight for researchers**

दिल्ली विश्वविद्यालय
University of Delhi

1, 2, 3 Data provided by the University of Delhi.

Published by

FRONTIER
ENTERPRISE

In partnership with

Lenovo + intel

Jicara Media
135 Middle Road, Bylands
Building, #05-03,
Singapore 188975

+65 9026 0649
+65 6815 2996

contact@jicaramedia.com
www.jicaramedia.com