Clalit accelerates research and innovation with AI and GPUs

The patient lies at the center of everything Clalit Health Services (Clalit) does. Founded in 1911, it provides healthcare services for half of Israel's population. Its 14 hospitals include 8 general hospitals, 2 mental health hospitals, 2 geriatric hospitals, and a children's hospital. It also operates community clinics, dental clinics, imaging facilities, and a lifestyle program. Clalit recently established an advanced AI platform based on Red Hat OpenShift AI and Red Hat® OpenShift®. It is leveraging this platform to build learning processes and algorithms to identify new trends, patient and disease behavior patterns, and more. The Red Hat technology allows it to rapidly build security capabilities users can use to address the concerns faced by their business.

Question: What makes Clalit unique among its peers?

Eyal Dviri, Innovation Team Leader in the Data Department at Clalit: Clalit is the biggest health organization in Israel. And it's a public healthcare organization, making us the payer and the caregiver. We focus entirely on the patient, providing everything medical that a human being needs.

We provide health services for half the population of Israel, so around 5 million people. We also hold 360-degree health data from our patients dating back to 2000. That means we have more than 20 years of electronic medical records (EMR) with everything from text and imaging to genomics. We have the data every health researcher dreams about.

This is the uniqueness of Clalit. Our challenge is accessing the insights—the patterns of disease progression and patient behavior—held within that data. For that, we need a solution for a platform that will enable researchers and developers access to a secure data platform. We have started with establishing a secured environment and data pipeline and, after stabilizing the platform we wanted to advance and enable machine learning capabilities to create AI products and processes and ML/AI workloads need GPUs to function optimally.

Question: Describe the first steps on your AI journey.

Dviri: When we started advancing ML/AI three to four years ago, we had two options for deploying an AI solution: on-premises or the cloud. We decided we wanted both. We found a cloud solution first; then, we saw the costs. We quickly learned that the costs would be difficult to control in a large organization like ours. So, we decided that on-premises would be better and bought an NVIDIA server.

We explored leveraging the NVIDIA GPUs with vanilla Kubernetes but quickly realized that this approach would not be sustainable. We have around twenty researchers using AI today, but that number will grow. We will eventually need to provide a robust AI platform for many more.

Question: What led you to Red Hat OpenShift AI?

Dviri: Our central IT department implemented Red Hat OpenShift to provide a modern solution for a wide range of use cases across the caregiving part of our organization. When they made it available to us, we adopted it immediately for our research needs. We didn't think twice. To address our ML/AI use cases, because we had OpenShift already, we decided to go with Red Hat OpenShift AI.

We needed a GPU solution and saw that Red Hat OpenShift AI was a promising solution. It provided many tools to make it easier to manage resources from the underlying Kubernetes platform. It would give us more control by providing a dashboard to see who was using what and how they were using it.

That was two and a half years ago. Our first research department projects using Red Hat OpenShift AI went live in August 2023. From now on, all research and machine learning projects will run on Red Hat OpenShift AI; every use case that requires a GPU processor will run on Red Hat OpenShift AI.

Question: Tell us about your first Red Hat OpenShift AI use cases.

Dviri: Our first use case is a joint venture with Harvard University Medical School that started before the LLM (large language model) OpenAI revolution. Two researchers from Harvard are taking different approaches to preparing data for the medical data foundation for future AI. Red Hat OpenShift AI is helping them process their large quantity of structured data using an NVIDIA GPU.

Our second use case uses LLM to pinpoint patients for preventive medication or closer inspections. We need to look at doctors' notes to find these people; for that, we use LLMs such as LLaMa. We upload the data—the text—and our LLM machine analyzes it. We are training it to make it more accurate. Then, Red Hat OpenShift AI allows us to deploy containers to download models and to use the GPU to analyze the data. We are still exploring LLMs, but it looks highly promising.

Question: What are your other Red Hat OpenShift AI use cases?

Dviri: Another use case allows us to develop image-based machine learning processes. It will enable us to search the DICOM format imaging, including CT and MRI scans and X-rays. Again, we need NVIDIA's GPUs to search DICOM format images, and Red Hat OpenShift AI allows us to do that.

My final use case—our 'supermodel'—uses Red Hat OpenShift and is already in production. We call it our 'supermodel' because it gives doctors at our hospitals input from our other hospitals. They can send a patient's ID to an OpenShift container which then queries our data warehouse and gives them more insight into the approach they're planning: whether it is safe to proceed or whether they should look deeper.

Question: How has OpenShift AI helped your researchers and data scientists?

Dviri: One of the most significant benefits is that it allows them to use the computing power they need in a scalable and comfortable manner. There is a good user interface that will enable them to work in a way they're used to working. It's a comfortable environment with plenty of support because they are a top priority for our Red Hat teams and us.

Red Hat OpenShift AI ensures they have the computing power they need to search text, search images, train models, and, in the future, process genomic data. We have found OpenShift AI to be beneficial in different ways, such as self-service to our data scientists and agility in the deployment of containers. To sum things up, the combination of OpenShift, OpenShift AI, and NVIDIA GPU, made the data scientists happy and they are looking forward to the expansion and newer versions of OpenShift AI solution.

Question: What achievements are you most proud of?

Dviri: I'm proud we can deliver secure research environments fast with Red Hat OpenShift AI. In April 2024 I was told our board was visiting our Innovation Center. Our data scientists wanted to demonstrate LLMs, and the team was asked to give them everything they needed. We did it. We gave the data scientists all they needed in two weeks, including building the notebook and downloading LLM packages from Hugging Face repository, a company that develops computation tools for building applications using machine learning. It was very swift.

As another example, two weeks before my vacation last July, I received a call saying Harvard needed OpenShift AI in two weeks. We needed to create an environment from scratch. But we had it ready within two weeks—two days before I went on vacation. Red Hat Consulting helped us do that. They opened a WhatsApp group with all the Red Hat people who would help, including those from overseas. The time-to-market was 'wow,' very fast and from nothing.

Question: Tell us more about how Red Hat has supported you.

Dviri: The team from Red Hat Consulting has been a great support. Without them, there would be no project; they're a critical part of the solution. We meet every week. They take the lead in sharing how best we can improve and scale our Red Hat OpenShift AI environment. It's an ongoing process. We're learning from them and becoming less reliant on them.

Red Hat supports us with solutions that address our customers' challenges. For example, a data scientist had a container for a MlFlow (machine learning flow), but it didn't work. The Red Hat team deployed a container image and configured it. When they gave it to the data scientist to test, they couldn't use it because of inner restrictions. So, the Red Hat team changed the permissions to what was needed and later even installed a different version. They kept talking until it worked.

Question: What's next for Clalit?

Dviri: We have two challenges ahead. The first thing is the increase in demand when people hear about our OpenShift AI platform. We need to figure out how to use our computing resources—our GPUs—optimally. We need to be able to give newcomers as many computing resources as they need while reserving enough for others. So, we're looking into best practices for onboarding new customers. We need to write a playbook detailing the steps they need to take.

The second thing is the pipeline. We need to understand best practices for deploying AI algorithms from training to production. Take the supermodel as an example. After we've trained the model in the research environment, how do we use OpenShift Pipelines to deploy it in a production environment? We must ensure the test and production environments are updated and in sync.

Containers are a critical part of our future. We have a lot of microservices that we can combine in containers to enable our doctors to help our patients. New research. New treatments. It's a win-win. We treat our patients better, and they spend less time in hospital. This is our motivation, and containers and OpenShift are shaping it. It's scalable because of OpenShift.

About Clalit Health Services (Clalit)

For more than one hundred years, Clalit Health Services has been Israel's leading and largest healthcare organization. Today, it is a non-profit integrated care organization for over four million patients, more than half of the country's population.

About Red Hat

Red Hat is the world's leading provider of enterprise open source software solutions, using a community-powered approach to deliver reliable and high-performing Linux, hybrid cloud, container, and Kubernetes technologies. Red Hat helps customers develop cloud-native applications, integrate existing and new IT applications, and automate and manage complex environments. A trusted adviser to the Fortune 500, Red Hat provides award-winning support, training, and consulting services that bring the benefits of open innovation to any industry. Red Hat is a connective hub in a global network of enterprises, partners, and communities, helping organizations grow, transform, and prepare for the digital future.

North America	Europe, Middle East, and Africa	Asia Pacific	Latin America
1 888 REDHAT1 www.redhat.com	00800 7334 2835 europe@redhat.com	+65 6490 4200 apac@redhat.com	+54 11 4329 7300 info-latam@redhat.com

facebook.com/redhatinc
twitter.com/RedHat
linkedin.com/company/red-hat

redhat.com