

Building self-learning AI

Aleph Alpha builds self-learning artificial intelligence on OCI

Based in Germany, Aleph Alpha is an independent artificial intelligence R&D startup. Their approach to AI represents a shift from “supervised” systems taught to complete tasks, such as identifying cars and pedestrians through labeled examples, to “self-supervised learning” networks that can find hidden patterns in data without instruction.

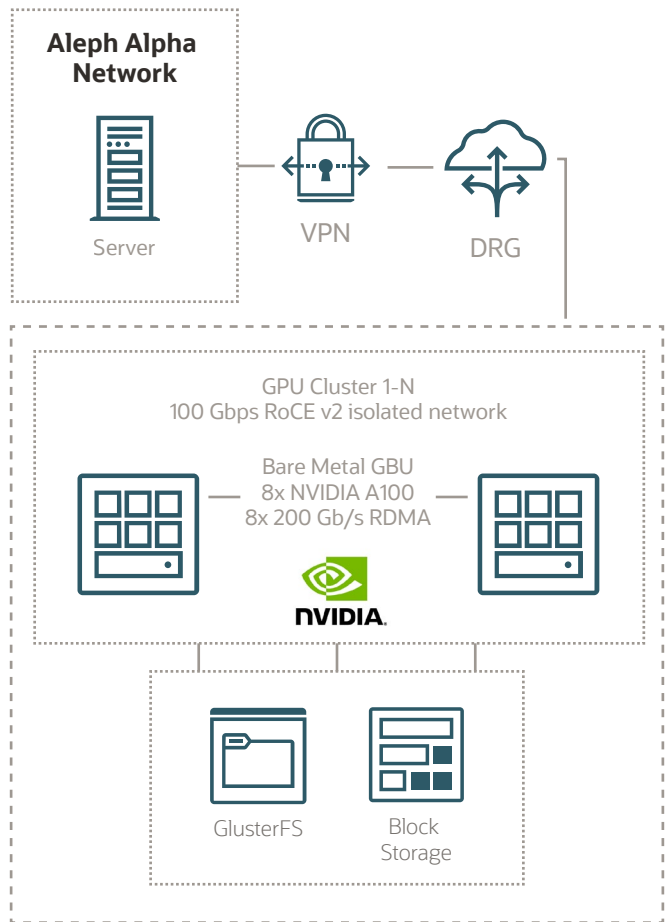
Enter OCI

Aleph Alpha, which recently closed a \$27 million dollar funding round in July 2021, is training a 13 billion parameter AI model on OCI, using hundreds of NVIDIA’s most powerful GPUs connected by high-speed networking. A second model holds over 200 billion parameters.

The new “self-learning” networks are able to answer questions never seen, match languages to cultures, and pepper answers with backup facts. These improvements are beyond the capabilities of the statistical prediction engines that are a standard of artificial intelligence today.

Self-supervised systems turn traditional software development on its head. Instead of tackling a specific problem in a narrow field, the new AI architects first build their self-learning models, let them ingest content from the internet and private datasets, and then discover what problems to solve. Aleph Alpha is fine-tuning its general AI model with specialized data in fields such as finance, automotive, agriculture, and pharmaceuticals.

The model that Aleph Alpha has created is unique in comparison to other existing AI models. Their training model is significantly larger and thus one GPU or one node may not be sufficient. When someone runs the training, they all reduce. This kind of optimization requires precise communication between nodes, so with



OCI’s engineering help, AA was able to deploy with A100 GPUs that are connected via RDMA in 8x 200 Gb/s connections.

Aleph Alpha also struggled with storage requirements. For their multimodal model, they had about 200 terabytes of super small, low-resolution images.

With OCI, the company was able to spread their model out over hundreds of high-speed connected GPUs and capitalize on Oracle’s advanced storage for CPU, RAM, and GPU.

“This is a new generation model, and in order to train those you need a new generation of hardware—the old GPU clusters aren’t sufficient anymore. On the industry side we have raised a lot of capital and partnered with Oracle. We’re building a way to translate an impressive playground task into an enterprise application that creates value.”

— Jonas Andrulis, Founder & CEO, Aleph Alpha

Learn more



Explore OCI’s HPC capabilities



Watch the video