

GPU Accelerated AI (Artificial Intelligence) Cloud for A Data Pro

*Machine learning enabled by a
blazingly fast storage solution*

About A Data Pro



A Data Pro provides content, data and business intelligence services. Enterprises around the world use A Data Pro to power their editorial, marketing, PR and risk & compliance efforts and products.

A Data Pro is a global company established in 1996 and is one of the leaders in content, data and business intelligence services. With a proven track record and 20+ years in business it has customers in many industries, in 80+ countries and more than 400 employees.

The Challenge

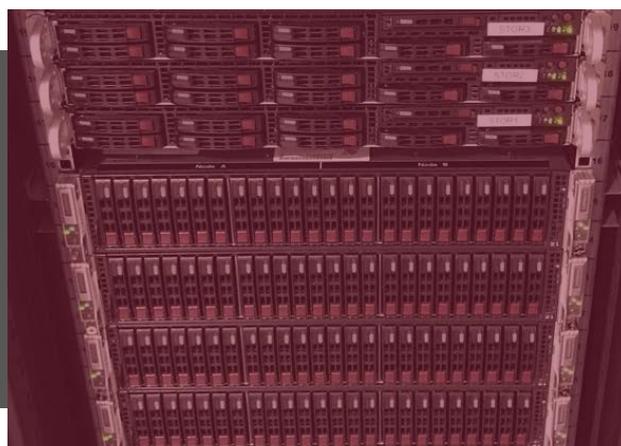


As a new line of business, A Data Pro developed a GPU-powered AI (Artificial Intelligence) which does semantic content processing of large volumes of data and can then analyse the sentiment of the content which is fed to the system. The solution uses machine learning, based on regular feedback loops and crunches extremely large sets of data.

For this project to work, A Data Pro needed to build a custom private cloud with enough processing power to perform the calculations needed to teach the AI. After extensive research on the options for building an extremely fast private cloud, A Data Pro chose to use a distributed, software-defined storage solution, for the speed, scalability and reliability a distributed storage system can provide.

Top Criteria

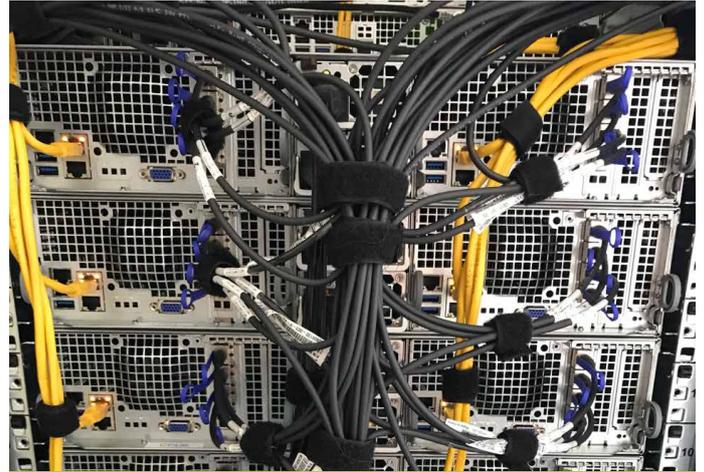
Among the most important factors were the fast random data access and low storage latency. In addition, the storage solution needed to provide scalability for storing rapidly growing amounts of data and parallelism in data processing, so to accelerate the machine-learning loops. A Data Pro was searching also for an easy-to-deploy storage solution with seamless integration with OpenNebula, which was the selected cloud management system (CMS).



The Solution

By leveraging the flexibility of StorPool's distributed SDS, the customer designed an optimized private cloud tailored to the initial needs of the project. Initially A Data Pro deployed StorPool Storage on three stand-alone storage nodes, using entirely SSD and NVMe drives connected to compute hypervisors by a redundant standard Ethernet network. The solution is running using KVM as hypervisor and 2 flavours of operating systems - Debian and CentOS for the compute and storage, respectively.

For the orchestration of the of the entire private cloud A Data Pro chose OpenNebula cloud management system. It was selected after test of several cloud management systems. OpenNebula turned out to be the prime choice - simple, yet powerful, reliable and fast.



StorPool was the only storage solution to meet and exceed all criteria of A Data Pro!

Key Benefits

- blazing speed
- scalability
- 24/7 support
- proactive monitoring
- complete solution out of the box

“StorPool offers a complete solution with exceptional performance and support. Anytime you need them, they are ready to help. StorPool is also doing constant monitoring on the storage and hosts and takes care of the health and uptime of our system. In this way they save us time and we can focus on our core activities. We are very happy with StorPool storage!”, said
Vasil Shivachev, part of A Data Pro.

Technical Details



Hardware Stack:

Intel CPU platform
Nvidia Tesla Graphic Cards
All-SSD Set-up
Redundant Mellanox network

Software Stack:

Linux
KVM
StorPool Storage
OpenNebula