NIS STC

NIS Group is one of the largest vertically integrated energy companies in Southeast Europe whose main activities consist of research, production and oil and natural gas refining, sales and distribution of a wide assortment of petroleum products, as well as the realization of projects in the fields of petrochemistry and energetics.

Challenges:

The main issue with the NIS STC project was coming up with an efficient way of storing and analyzing data that makes images of underground structures and delivering those images to scientists who will further analyze them.

The Solution:

We are creating a global solution by integrating different standalone technologies into one system. The system consists of distributed storage technologies integrated with data analysis technologies.

The Result:

The result is a financially efficient way to make their operations much faster, data transfer more reliable and secure, and deliver results to scientists wherever they are.

"Quick, intelligent and reliable, turnkey like approach to provisioning custom IT infrastructures"

Leonid Stulov
CEO NIS STC

"Braineering solution delivery always has "end to end" aspect to it. They deliver what you want, and more because they know what you need before you do."

NIS STC

Company Background:

NIS Group is one of the largest vertically integrated energy companies in Southeast Europe. They have their R&D center called NIS STC (The Scientific and Technological Centre). STC and its qualified and experienced personnel are in charge of the preparation of geological-exploration projects, and the estimation of reserves, the processing of seismic data, developing databases and software, as well as well-equipping designs. The central laboratory, which facilitates the preparation of geological-exploration projects through analyses (core analyses, oil, and fluid characteristics), is located in the STC. It practically means, among other things, that they analyze vast areas of underground structures to predict where to extract oil. The data gathered is further analyzed by the software that can render that data into pictures and charts that humans can understand. This way, the oil company knows where should be the most significant probability of successful oil extraction.

The Challenge:

The main issue with the NIS STC project was coming up with an efficient way of storing and analyzing data that makes images of underground structures and delivering those images to scientists who will further analyze them.

The Solution:

We are creating a global solution by integrating different standalone technologies into one system. The system consists of distributed storage technologies integrated with data analysis technologies. For this system to work, we needed to make a big powerful computer. This computer consists of:

- Data Storage
- Computing power
- Network: that assures interconnection between different system elements
- Technologies for analyzing and delivering results of computing to humans.

This computer is much like any other. The only difference is the scale of power, storage, and speed. In other words, this computer is so big that you cannot place it into the classic computer case. You instead need a building called the data center.

Data Storage:

We needed storage that can protect that data and make it highly available. We used distributed storage technology, which means that different data sets are spread across other discs, magnet, or flash storage devices. Storage virtualization technologies allow the software to see the data distributed across multiple physical locations, as they were on a single storage pool. The system is not being aware of distribution at all. With this technology, you can combine different storage devices (tape, disc, flash) and make software see it as one pool. The main advantage of this technology is the ability to tier data based on the data value and put less valuable data on cheaper hardware, and more important and expensive data on more costly but more reliable hardware. It also stands for the usefulness of data. Hot data is data that needs to be accessed right away, it is usually business-critical, and you cannot afford any latency, so the system puts it in faster flash devices. If the data is no longer in active use and might not be needed for months, years, decades, or maybe ever, it is considered cold, and the system puts it in slower but cheaper devices, with higher latency. The data storage system also has advanced deduplication, compression, compaction, snapshot, and thin provisioning features. All these features elevate the security and availability of the data and make it more cost-effective.

Computing power:

Computing power for this project was already provisioned by the existing HPC (High-Performance Computing) system. Our main goal was to make the storage system available to the HPC system by integrating it through a high bandwidth network.

Network:

The high bandwidth network consists of two physical switches and many virtual ones that connect the system through interfaces that can transfer data as fast as 100Gbps. The network also has advanced zero trust security features based on OSI models, for further protection of data in transit.

Analyzing and delivering:

This part consists of VDI technology. VDI stands for Virtual Desktop Infrastructure technology. It is, essentially, a mobility technology that enables the virtualization of desktops. In a traditional desktop PC setting, from a human perspective, your desktop is GUI (Graphic User Interface) that helps you control your apps and do your business. It is delivered to you by a cable, which connects your computer case to a monitor device that displays that GUI for you. However, if we look at that cable as wireless WAN, that delivers your desktop to your mobile device or laptop instead of your monitor, you get VDI. This technology adds a layer of abstraction to physical equipment and enables the use of that equipment from a remote location. Using high power GPU (Graphics Processing Unit, or simply graphic card) and leveraging VDI technologies, we can transfer and analyze the data. That data is on the storage system that we built and on the preexisting HPC system, and deliver it to scientists for further analysis.

The Result:

The result is a financially efficient way to make their operations much faster, data transfer more reliable and secure, and deliver results to scientists wherever they are.

1. The system performs faster than ever before.
2. It is more reliable with no data loss in transit, and more resilient to catastrophic events through a redundant array of independent discs technologies.
3. The data is much more secure and available
4. Scientists are more flexible than ever because we deliver results directly to them wherever they are.