## Northern Data takes HPC to a new level of cost and energy efficiency with AMD technology

AMD EPYC<sup>™</sup> CPUs and AMD Instinct<sup>™</sup> GPUs enable more scalable and power efficient processing to bring AI or ML to more customers.

### CUSTOMER



### INDUSTRY

HPC infrastructure provider

## CHALLENGES

Provide HPC performance for AI and ML at a more affordable price

## SOLUTION

Deploy AMD EPYC<sup>™</sup> CPU-powered GIGABYTE G292-Z20 servers with AMD Instinct<sup>™</sup> GPUs

#### RESULTS

Half the cost per ML training cycle and 30-40 percent lower power consumption

## **AMD TECHNOLOGY AT A GLANCE**

AMD EPYC 7402P CPUs with 24 cores AMD Instinct MI50 GPUs

**TECHNOLOGY PARTNER** 

## **GIGABYTE**<sup>®</sup>

The high-performance data center market has been rapidly expanding and U.S. brands have tended to be the predominate providers. But Northern Data is going to change that. The multi-billion-dollar European tech provider is aiming to offer HPC cloud services such as machine learning (ML) to a wider audience. Designed for HPC workloads and modern software architectures, Northern Data's data centers are to more than 90 percent powered by green energy.

To achieve this goal, Northern<br/>Data relies on technology<br/>partners that could deliver the<br/>performance required with the<br/>optimal cost efficiency.<br/>Together with GIGABYTE,<br/>Northern Data found that AMD<br/>EPYC and AMD Instinct<br/>technologies delivered the<br/>scale and affordability that the<br/>company's customers demanded. Northern Data"We manage<br/>power consu<br/>30 to 40 p<br/>oother cloud<br/>Michel Bout<br/>Softwa

is planning to provide a local service compliant with recent GDPR laws, which is an increasing concern for European customers.

# Faster, scalable and more affordable ML training

"We are going to obtain some market share from the hyper scalers for AI, Rendering, and other HPC applications," explains Michel Boutouil, General Manager of Northern Data Software GmbH. "We bring affordable HPC applications to the market, because HPC is still expensive for several reasons, which prevents small and medium-sized enterprises from entering this field." This led Northern Data to both AMD EPYC processors and AMD Instinct GPUs, which together delivered excellent scalability and lower power consumption.

"AMD processors are of high performance and power efficiency," says Boutouil. "When we evaluated the AMD Instinct MI50s, they performed particularly well in large clusters, and their power consumption is also very low." Northern Data also appreciated the direct dialogue with AMD engineers to optimize the AMD Instinct GPUs' performance. "We are in close collaboration with the AMD team. We are confident the AMD product will be a great choice for our customers for ML applications, especially at large scales."

To confirm the performance of the AMD Instinct GPUs, Northern Data performed a series of tests, starting with a TensorFlow implementation

aimed at training ML frameworks.

Northern Data measured how

many seconds an AMD Instinct

MI50 GPU took to train ML with

300,000 images compared to a

typical instance from a major cloud

"We managed to lower the power consumption about 30 to 40 percent for a comparable workload versus other cloud platforms."

Michel Boutouil, General Manager of Northern Data Software GmbH

l, General hern Data mbH seconds. Northern Data stated that this would considerably lower the cost of training ML systems. A full training cycle could be much more cost efficient than before. "On the compared instance, for example, it would cost

compared instance, for example, it would cost €1,987 (\$2,339) to run one experiment," says Loesch. "We can do it with AMD solutions for €1,100 (\$1,295). So that's almost half the price."

This could be the difference between employing ML and it not being cost effective. "Typically, in machine learning, you would run several experiments with each team deploying several experiments a day," says Loesch. "You want to have a scalable solution. If you are spending two thousand Euros on an experiment when you do 20 a day, it can get quite expensive very fast.

## Linear scaling with AMD Instinct GPUs

More and more companies are now using Al and ML like this on a regular basis. "Social media will be a big one," says Boutouil. "For example, companies that deploy these sorts of technologies would use this for research. Natural Language Processing (NLP) is also quite big, which we're going to evaluate in the future as well."



Northern Data expects its AMD technology to deliver these capabilities to a

much wider range of clients. "With NLP becoming more and more expensive, we could give AI startups a chance to be competitive against the bigger players, which have millions to blow on training."

Further underlining the possibilities, Northern Data found AMD Instinct GPUs scaled in a brilliantly linear fashion, with TensorFlow ResNet50 on AMD Instinct MI50 GPUs improving from around 400 images per second with one

GPU to around 2,250 images per second with eight – almost exactly the eight times performance you would hope for. Similar scalability was shown when running a Blender<sup>®</sup> 3D render in a virtual machine on the AMD Instinct MI50 GPUs.

Even better was the power efficiency, which more than doubled the amount of work completed per watt with eight GPUs compared to one. This is because eight GPUs can be installed on a single AMD EPYC CPUpowered server without losing any bandwidth, thanks to the 128 PCIe Express® 4 lanes per single socket and up to 160 per dual socket. In order to fulfill this high demand workload, the GIGABYTE G292-Z20 was chosen for its efficient design and topology that allowed for maximum throughput for the AMD Instinct MI50. For a 2U GPU-dense chassis, thermals can be a challenge, but this server excelled without throttling performance. "For distributing training workloads, the extremely large bandwidth towards the GPUs really helps," says Loesch. These positive results led Northern Data to place a huge order for 4,366 single-socket GIGABYTE servers powered by AMD EPYC 7402P processors, each one equipped with eight AMD Instinct MI50 GPUs, for a grand total of 34,928 accelerators. Of these servers, around 2,000 have been deployed so far.

### Low costs and more speed for new industries

Northern Data is already seeing its investment deliver outstanding results. "We managed to lower the power consumption about 30 to 40 percent for a comparable workload versus other cloud platforms," says Boutouil. "The reduction in carbon footprint is especially important to us.

""For distributing training workloads, the extremely large bandwidth towards the GPUs really helps."

Ferdinand Loesch, Software Engineer at Northern Data Furthermore, the low costs empower more people, who previously weren't really capable of utilizing the technology, due to economic factors. We can deploy amazingly fast data centers, and we are able to scale quickly to run large GPU clusters. This combination also leads us into new industries. For example, the healthcare industry, biotech or MedTech companies, which have a need for strong GDPR assurances. The carbon footprint is especially important to us, which the lower power consumption really helps."

"We see our data centers not as a customer data center, we see IT as our customer, and we try to build a perfect setup that works together,"

In addition, we really like the flexibility of the open-source approach.

"We like the service AMD provided, as well as the flexibility, the open-source approach, the way of thinking, and of course, the excellent and costeffective technology. That is why we chose AMD."

> Michel Boutouil, General Manager of Northern Data Software GmbH

concludes Boutouil. "We worked with colleagues from GIGABYTE and AMD to find the optimal setup for the servers, and the optimal density because space is money, and we buy as little space as we need. We like the service AMD provides, as well as the flexibility, the open-source approach, the way of thinking, and of course, the excellent and cost-effective technology. That is why we chose AMD."

## WANT TO LEARN HOW AMD EPYC<sup>™</sup> PROCESSORS MIGHT WORK FOR YOU?

Sign up to receive our data center content amd.com/epycsignup





## **About Northern Data**

Northern Data AG develops and operates global infrastructure solutions in the field of High-Performance Computing (HPC). With its customer-specific solutions, the company provides the infrastructure for various HPC applications in areas such as bitcoin mining, blockchain, artificial intelligence, big data analytics, IoT or rendering. The internationally active company is today a leading provider of HPC solutions worldwide. Northern Data offers its HPC solutions both in large, stationary data centers and in mobile high-tech data centers that can be set up at any location worldwide. The company combines self-developed software and hardware with intelligent concepts for a sustainable energy supply. The Northern Data group currently employs around 200 people. For more information visit <u>northerndata.de</u>.

## **About GIGABYTE**

GIGABYTE offers a comprehensive product lineup that aims to "Upgrade Your Life." With expertise encompassing consumer, business, gaming, and cloud systems, GIGABYTE established its reputation as a leader in the industry with awardwinning products including motherboards, graphics cards, laptops, mini PCs, and other PC components and accessories. As a perennial developer in PC and server hardware and solutions, GIGABYTE is poised to extend its scope in business servers and cloud systems with hardware and software solutions that integrate AI and AloT applications to allow customers to capture, analyze, and transform digital information into economic data, accelerating businesses' success from edge to cloud. For more information visit gigabyte.com.

## **About AMD**

For over 50 years AMD has driven innovation in high-performance computing, graphics, and visualization technologies the building blocks for gaming, immersive platforms, and the data center. Hundreds of millions of consumers, leading Fortune 500 businesses, and cutting-edge scientific research facilities around the world rely on AMD technology daily to improve how they live, work, and play. AMD employees around the world are focused on building great products that push the boundaries of what is possible. For more information about how AMD is enabling today and inspiring tomorrow, visit <u>amd.com/EPYC</u>.

All performance and cost savings claims are provided by Northern Data and have not been independently verified by AMD. Performance and cost benefits are impacted by a variety of variables. Results herein are specific to Northern Data and may not be typical. GD-181

©2021 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, EPYC, AMD Instinct, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Blender is a registered trademark of the Blender Foundation in EU and USA. PCIe is a registered trademark of PCI-SIG Corporation. TensorFlow, the TensorFlow logo and any related marks are trademarks of Google Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

