

Accelerate AI-Driven Results

How Pure Storage is Helping Customers



Introduction

Organizations that are more mature in their adoption of analytics and AI technologies have a competitive advantage, with AI is transforming business, government, and scientific research everywhere. But data doesn't analyze itself, and legacy systems often can't support the massive data pipelines required for modern analytics, machine learning, and other advanced AI techniques. Pure Storage® enables advanced AI across different industries with a highly scalable and parallel platform that delivers ultra-fast, all-flash performance that accelerates insight and results for data scientists and researchers everywhere.

Pure Storage FlashBlade® is enterprise storage built for AI, which delivers multi-dimensional performance via a scale-out storage platform that consolidates data for AI workloads. FlashBlade is also a key component of an NVIDIA and Pure AI Ready Infrastructure (AIRI®), providing end-users with a turnkey compute, storage, and networking solution that is simple to manage; can start small and scale large; and be agile and evergreen, ready for future AI innovations.

Here are just a few examples of customers across different industries that are leveraging powerful solutions from Pure Storage to help with the ever-increasing demands of their AI technology.

STORY 1 Crater Labs	03
STORY 2 GraphCore	05
STORY 3 Gwangju Institute of Science and Technology (GIST)	07
STORY 4 Kakao	09
STORY 5 NCSOFT	11
STORY 6 Olympus	13
Conclusion	15





Crater Labs Aims High and Fuels its Journey with Pure Storage

In 2017, Alexei Gavriline and Khalid Eidoo co-founded Crater Labs with a lofty goal—to help companies tackle their most ambitious, groundbreaking projects. With world-class AI and ML researchers and a high-performance computing infrastructure, Crater Labs makes it possible and profitable for clients to achieve their goals. Data is at the center of every Crater Labs project.

The company's researchers must be able to access, analyze, and move data at will as they run their experiments. Having maxed out capacity and performance on both cloud and direct attached storage, Crater Labs turned to Pure Storage.



“Pure gives us credibility as we take on transformative AI initiatives for clients. It proves we understand what it takes to run high-risk, high-reward research projects successfully and cost-effectively.”

ALEXEI GAVRILINE,
CO-FOUNDER & PRESIDENT,
CRATER LABS



Fueling AI/ML Research with Big Data

Crater Labs AI and ML solutions are making self-driving cars better at detecting cyclists and pedestrians; preventing fraud by spotting anomalies among billions of financial transactions; and sparking innovation in logistics, healthcare, telecom, and beyond. Its researchers are central to this effort.

“We invest a lot in world-class research and development talent, and we need to make sure they can run their experiments without interruption or delay,” says Alexei Gavriline, president of Crater Labs. “The faster they work, the faster we can move the needle for our clients.” The company needed high-performing infrastructure, but their direct attached storage was maxed out and couldn’t keep up with the data-intensive workloads.

“We were burning out hard drives left and right because of the high I/O requirements of our applications,” says CTO Khalid Eidoo. “We needed the flexibility and performance to handle parallel workloads across our hybrid cloud environment.”

Delivering Transformative Results to Clients Faster

By adopting Pure Storage FlashBlade, Crater Labs combines the flexibility and abstraction of cloud storage with on-premises control and performance. That makes it easy to collaborate with clients on large-scale AI projects.

“Using [Amazon] S3-compliant data stores on FlashBlade, we can easily move data back and forth, accessing our clients’ files with very low latency,” says Eidoo. “Plus, initiating new experiments is quick—we simply spin up a container and point the volume to the FlashBlade. No more data shuffling.”

Faster performance means Crater Labs can run multiple experiments simultaneously, allowing researchers to work unimpeded by resource constraints and deliver results faster. And researchers don’t have to worry about losing data, as snapshots can be restored in minutes.

Using FlashBlade, research teams have helped clients analyze SEC data in three days instead of 10, calculate delivery routes for thousands of trucks in two-thirds less time, and detect billing anomalies with up to 93% accuracy for telco and utility companies.

Crater Labs Moves Ahead

Crater Labs looks to scale and expand its business globally with help from Pure. “Pure gives us credibility as we take on transformative AI initiatives for clients,” says Gavriline. “It proves we understand what it takes to run high-risk, high-reward research projects successfully and cost-effectively.”

About Customer

Toronto-based Crater Labs takes on high-risk, high-reward AI/ML projects to create lasting competitive advantage for its clients, providing the research acumen and infrastructure to do it cost-effectively.

www.craterlabs.io



Geo

North America



Industry

High Tech


Results

- Accelerates research by running parallel workloads simultaneously
- Easier to transfer and work with client data using Amazon S3 data stores
- Simplifies storage management to spin up experiments quicker

Impact on Crater Labs

- Streamlines collaboration with clients on groundbreaking AI projects
- Allows researchers to run experiments unimpeded for faster result
- Supports a hybrid cloud strategy for both performance and flexibility





Graphcore Breaks the Barriers on AI Exploration

Graphcore's IPUs are disrupting AI, setting the global standard for machine intelligence compute and creating opportunities across everything from drug discovery to decarbonization.

The IPU has a unique architecture specifically designed for AI compute needs. It delivers a huge performance uplift and lets researchers explore new types of work that were previously not possible or functionally challenging using legacy processor architectures such as the GPU.

As a key part of the end-to-end AI optimized technology stack, Pure Storage supports Graphcore's goals to break through the shackles placed on AI exploration.

GRAPHCORE

"Getting the best from our IPUs demands storage with high performance, low latency, high throughput, and distributed scalability. FlashBlade has delivered on all these fronts."

STUART CORNELL,
SYSTEMS OPERATIONS
MANAGER, GRAPHCORE



Smoothing Access to Data—Innovating With AI at Speed and Scale

The challenge with AI models is that they require a lot of data to train, improve, and refine them. Having the right storage and right configurations is critical to creating an efficient and reliable infrastructure for AI workloads.

Stuart Cornell, Systems Operations Manager at Graphcore, explains, “Graphcore systems deliver extremely high-performance AI compute. To get the most out of our IPU processor, it is important to have storage systems that also meet the unique and demanding requirements of AI, and Pure delivers that.”

Delivering a Flawless Storage Experience

Graphcore deploys Pure in Graphcloud — its cloud-based machine-learning (ML) platform that helps customers expand from experimentation and proof of concept projects to production systems. Graphcloud runs on large IPU Pod systems, powered by several hundred IPU processors, which all rely on Pure FlashBlade®. The all-flash unified fast file and object (UFFO) platform enables the company to consolidate data across Graphcloud’s machine learning operations (MLOps) pipeline to keep the IPU Pods fed.

“Getting the best from our IPUs demands storage with high performance, low latency, high throughput, and distributed scalability,” says Cornell. “FlashBlade has delivered on all these fronts. And it’s simple. In the best possible way, Pure means I can forget about the storage. This flawless experience also reflects what we want for our customers.”

FlashBlade also brings significant power consumption efficiencies to the overall system. FlashBlade has high performance density, delivering an enormous amount of throughput for a small unit. Given Graphcore IPUs are throughput hungry, Pure can replace their legacy storage equipment that would inevitably take up much more space and power.

Engineering for the Next Generation of Compute

Cornell concludes: “Pure is engineering the type of storage technology that AI needs. Graphcore’s adoption of Pure brings intelligent compute and storage together to create a converged infrastructure solution capable of serving AI and ML workloads at any scale.”

About Customer

Intelligence Processing Unit (IPU) inventor Graphcore is setting new standards for AI compute and helping innovators make breakthroughs in machine intelligence that positively impact everything from drug discovery to decarbonization.

www.graphcore.ai



Geo
EMEA



Industry
Technology

Results

- Delivered high bandwidth, throughput, low latency for AI compute
- Simple installation, setup allows IT to set and forget storage
- Flexibility to support variety of data types on one platform

Impact on Graphcore

- Scale on the fly to support new and growing customer demand
- Focus on developing world leading AI accelerators
- Rapidly expand presence in new industries





FlashBlade Optimizes GIST's AI Infrastructure

As a global institution leading education and research in science and technology, GIST sought to improve its AI computing facility so that students can conduct research more conveniently. Needing parallel infrastructure with proven results in numerous research cases, GIST chose Pure Storage's AI-ready architecture and storage.

With AIRI powered by FlashBlade and NVIDIA's DGX A100, GIST enhanced its AI infrastructure and allowed students to conduct their research more efficiently and reliably. Pure1® also allowed the IT team to monitor storage anytime, anywhere, increasing the efficiency of infrastructure management.



“You need to consider storage to drive the success of AI projects. FlashBlade processes large amounts of data reliably at high speed, allowing us to build an infrastructure that always provides the optimal performance.”

JARGALSAIKHAN NARANTUYA,
PROFESSOR, AI GRADUATE SCHOOL,
GIST



GIST Improves AI Computing Environment with AIRI on FlashBlade

GIST is a research-intensive university established by the South Korean government to drive the nation's IT innovations, promote industry-academic collaborations, and train elites who will lead the 4th industrial revolution. In line with Gwangju City's plan to establish the nation's first AI industrial cluster town consisting of data centers and testbeds to incubate AI startups, The AI graduate school of GIST is striving to nurture professionals who will lead this integrated AI complex.

To train these future experts and expedite their research, GIST wanted to improve its AI computing facility. To do this, GIST needed parallel storage that could quickly process large amounts of data for AI/ML workloads. GIST also required one that could be easily operated by its small IT team, and one that has proven its performance in research areas including automobile, energy, and healthcare.

FlashBlade Offers High Performance that Empowers AI Researchers

This led GIST to augment its AI computing cluster center with AIRI, a full-stack AI-ready infrastructure powered by FlashBlade and NVIDIA's DGX A100.

Designed to deliver maximum performance for unstructured data, FlashBlade offers unified fast and object storage platform quickly delivered

massive datasets to AI/ML applications, boosting GPU utilization. In addition, when students uploaded ML training data to the cluster center, the system immediately delivered data to storage and DGX servers using the NFS protocol, which in turn accelerated research projects for developing AI in automobile, healthcare and energy.

Thanks to the reliability of FlashBlade and the real-time storage monitoring and predictive support of Pure1, the IT team was able to manage the system with ease. "You need to consider storage to drive the success of AI projects. FlashBlade processes large amounts of data reliably at high speed, allowing us to build an infrastructure that always provides the optimal performance," shared project lead, Professor Jargalsaikhan Narantuya.

GIST to Lead R&D with AI-optimized Computing Environment

GIST's AI/ML projects are growing daily, but FlashBlade offers unmatched performance and scalability that allows GIST to meet them all with confidence. Leveraging this modern infrastructure, GIST plans to continue research in various fields to educate AI leaders of tomorrow and develop technologies for the future.

About Customer

Gwangju Institute of Science and Technology is a renowned research-intensive university in Korea. GIST ranked no.4 in citations per faculty in 2021 QS World Ranking and no.1 in 2020 startup achievements in Korea.

www.gist.ac.kr



Geo
APJ



Industry
State and Local Government
and Education

Results

- Delivers data faster to AI/ ML applications, enhancing GPU utilization
- Able to monitor and manage storage capacity and performance with ease
- Supports NFS protocols, enabling swift delivery of data to storage and GPU

Impact on GIST

- Accelerates projects with fast and convenient AI computing environment
- Simplifies management, allowing the institute to focus on R&D
- Supports a wider range of projects with scale-out architecture





Kakao Optimizes AI Infrastructure with Pure Storage

Kakao operates Korea's No.1 instant messaging app, KakaoTalk, and provides services in diverse fields including entertainment, commerce, finance, and mobility to consumers around the world. Kakao is engaged in diverse AI projects to improve its services with data collected from these platforms and needed storage that could handle the massive influx of data from AI applications and deliver maximum performance from GPUs.

By deploying Pure Storage FlashBlade, Kakao freed up computing resources by eliminating network and storage bottlenecks, and significantly reduced the time and cost associated with operating and managing data.

kakao

“We needed fast and reliable storage for AI projects, and FlashBlade was an optimal choice for us. Since adopting FlashBlade, we are able to use data more efficiently and significantly reduce the cost of managing infrastructure.”

IL-SEOP JIN,
MANAGER,
INFRASTRUCTURE PLANNING,
KAKAO CORPORATION



Operating an AI Cloud Platform

With a mission of making a better world with people and technology, Kakao provides messaging apps and other communications services to consumers and businesses globally. Most recently, its COVID-19 vaccine booking system and multiple location-based services enabled people to quickly and easily find nearby locations to receive their vaccines.

With the onslaught of the pandemic, increase in telecommuting, and new service launches, Kakao needed a better infrastructure to handle the growing traffic and data on its services.

Kakao is also operating research teams to develop AI models for its diverse services and allocates computing and storage resources to these teams utilizing an on-premise storage platform. They required an extremely efficient storage solution supporting multiple servers. The storage also needed to be resilient in the event of data disasters, scalable to accommodate a large amount of data, and easy to operate. After evaluating solutions from several vendors, Kakao chose FlashBlade as their unified fast file and object storage platform because it demonstrated the highest performance with its massively parallel architecture.

Streamlined Computing Resources and Data Management

With FlashBlade, Kakao was able to eliminate network interface and storage bottlenecks, securing high-performance computing resources from GPUs.

“We needed fast and reliable storage for AI projects, and FlashBlade was an optimal choice for us. Since adopting FlashBlade, we are able to use data more efficiently and significantly reduce the cost of managing infrastructure,” shared Il-seop Jin, Manager of Infrastructure Planning team for Kakao.

In addition, Kakao gained better control over its data by utilizing a consolidated storage platform offered by FlashBlade. Seamlessly integrated with in-house control systems, FlashBlade provided intuitive management of data pipelines, eliminated data silos, and reduced operational and management costs through deduplication and compression. The RapidFile Toolkit provided by Pure Storage also dramatically reduced the time required for data querying and deleting unnecessary data.

Kakao to Provide Diverse AI-powered Services for Customers

As a result of the early success with FlashBlade, Kakao has been steadily increasing its reliance on FlashBlade over the past four years. With FlashBlade, Kakao plans to continually carry out IT innovations and AI projects, expanding its services to meet the diversified needs of its customers.

About Customer

Kakao Corporation is a ‘mobile life platform’ company dedicated to renewing daily lives by connecting people to people, and people to technology with flagship services including KakaoTalk, Melon, and Daum.

<https://www.kakaocorp.com>



Geo
APJ



Industry
IT, Platform

Results

- Secures computing resources from GPUs by eliminating bottlenecks
- Improves efficiency with central management, compression, and deduplication
- Scales infrastructure without additional management or operational burdens

Impact on Kakao

- Allows Kakao to improve the use of AI technology and data
- Rapid response to changing work environments and customer needs
- Boosts the quality of services and the ability to embark on new endeavor





NCSoft Boosts Data and Storage Efficiency for AI

NCSoft is a global game company that develops popular PC and mobile games including Lineage, Lineage Blade & Soul, LineageM, and Lineage2M. Recently, NCSoft began using AI to provide services that strengthen its development and service capabilities. To do this, NCSoft needed a storage solution to efficiently manage complex data and decided to implement a simple and AI-optimized solution, Pure Storage FlashBlade.

After deploying FlashBlade, NCSoft significantly reduced the time and value associated with storage and data management while gaining the flexibility to scale its on-premise infrastructure at no additional cost.



“FlashBlade is easy to deploy and operate. Both our developers and administrators are satisfied with the all-flash architecture’s performance.”

JIN-HYUK JANG,
DIRECTOR,
IT INFRASTRUCTURE
NCSoft



Accelerating AI Projects for Today's Gaming Experiences

NCSOFT is a leading game developer and publisher. Early on, NCSOFT began using AI technologies to develop ML tools to improve the quality of graphics and reduce repetitive and manual tasks in game planning, artmaking, and programming.

To further this effort, the company established an AI Center consisting of Game AI Lab, Speech Lab, Language AI Lab, Knowledge Lab as well as Vision TF Lab.

While carrying out these projects, NCSOFT had a hard time managing data. It took a long time for the company to move, load, and manage the data of each user, stage, and device. The company also had a hard time filtering unnecessary data, which led it to seek out a storage solution that could reduce the time of these tasks.

Expedites Timelines and Simplifies Deployment

Requiring high-performance storage for its bare metal, GPUvm and CPUvm environments, NCSOFT was interested in FlashBlade, the leading unified fast file and object storage platform (UFFO), for its simplicity, scale and high performance for running AI workloads. After completing a proof of concept (POC) and on-site testing, NCSOFT decided to implement FlashBlade, which reduced task times and improved efficiencies with easy configurations and management.

The company is now able to move and load data on the same day, which previously took several days to accomplish. The operations team can also quickly allocate volumes without having to assign personnel for the job. "FlashBlade is easy to deploy and operate. Both our developers and administrators are satisfied with the all-flash architecture's performance," shared Jin-hyuk Jang, Director of IT Infrastructure team for NCSOFT.

Furthermore, NCSOFT significantly reduced its total cost of ownership (TCO) by building an AI environment that is as scalable as the cloud while having it on-premise. FlashBlade offers native file and object architecture that allows NCSOFT to utilize S3 at no additional cost.

NCSOFT to Innovate with AI and Containers

NCSOFT is pursuing diverse innovations with AI to strengthen its leadership in the industry. To reach this goal, NCSOFT has been steadily increasing the scale of FlashBlade for several years with positive outcomes, and recently added new blades in bulk to accommodate more AI projects.

NCSOFT will continually use the AI and cloud-optimized performance of FlashBlade and is considering the Portworx industry-leading services for its container-based applications.

About Customer

NCSOFT is one of the leading online gaming publishers in the world, renowned for its technologies and services. Starting with the establishment of the U.S. subsidiary in 2000, NCSOFT is sharing the value of joy with more than 60 countries worldwide.

kr.ncsoft.com



Geo
APJ



Industry
Media, Entertainment

Results

- Facilitates data management and reduces the time needed for processing data
- Allocates volumes easily and scales with cloud like simplicity without draining workforce
- NFS, SMB and S3-compatible, replicating data to S3 at no additional cost

Impact on NCSOFT

- Reduces operational costs and facilitates the management of growing data
- Reduces time and effort with easy-to-deploy and operate infrastructure
- Accelerates AI projects to develop and improve diverse services





Supporting Olympus' AI Development with Pure Storage FlashBlade and AIRI

Olympus is a leading manufacturer of medical devices with an ambitious vision.

"In research and development related to AI, a large amount of image data is generated, and the infrastructure for storing that data becomes more important," says Mr. Shuta Yanagita, Software Strategy Manager, Software Strategy at Olympus.

"In addition, as the number of AI development programs increases, we need an environment that efficiently uses GPU resources. Creating a flexible AI development environment that solves these issues became a new challenge for the company."

OLYMPUS

"Pure Storage and AIRI prepare a platform for innovation to make advancements in AI development."

MR. SHUTA YANAGITA,
SOFTWARE STRATEGY MANAGER,
SOFTWARE STRATEGY, OLYMPUS



Improving R & D efficiency and strengthening competitiveness

Olympus has set out a policy of utilizing AI to improve customer value.

Based on this policy, the Software Strategy Group, which is in charge of building a software development environment, is working to improve the development environment with an emphasis on improving the efficiency of AI research and development, and strengthening the company's competitiveness. Not only did the infrastructure need to process immense amounts of data in various formats and at high speeds, it also needed to support rapid deployment in a container environment, powered by Red Hat OpenShift.

However, the company's existing infrastructure couldn't keep up with the growing volumes of imaging data and the fast pace of development. To find a solution, Software Strategy Group turned to Tokyo Electric Device (TED), a long-time partner who introduced Olympus to Pure Storage FlashArray™ years earlier.

Achieving a rapid AI development environment for developers

On TED's recommendation, Software Strategy Group adopted AIRI as a new AI development environment. Based on Pure Storage FlashBlade and NVIDIA

systems, the reference architecture for AI provides the performance and capacity required to run intelligent models on a large scale.

"Pure Storage and AIRI prepare a platform for innovation to make advancements in AI development," says Mr. Yanagita.

Using AIRI, with the speed and flexibility of FlashBlade, development teams can go to market faster with new AI solutions that have the potential to transform the market. And by powering their Kubernetes workloads with OpenShift and Pure, Olympus developers can provision all the resources needed to build, test, and deploy applications on demand.

"By running Red Hat OpenShift, developers can quickly spin up container-based development environments for themselves," says Mr. Yasuhiro Wada, Software Strategy U2, Software Strategy at Olympus. "Compared to building from a virtual machine on a runbook basis, we achieved a reduction in human error and a dramatic improvement in building speed."

Software Strategy Group sees enormous potential for growth, especially given the performance of FlashBlade. "Pure Storage FlashBlade is so reliable and stable," says Mr. Konno, Software Strategy U2, Software Strategy at Olympus.

About Customer

As a leading MedTech company, Olympus offers a wide range of products and services. We aim to contribute to improving the health and QOL of people around the world in diagnosis and minimally invasive treatment with innovative technology and manufacturing technology.

www.olympus-global.com



Geo
APJ



Industry
High Tech, Healthcare

Results

- Secured a stable AI environment with AIRI
- Achieved a self-service AI development environment
- Realized an efficient storage environment for container delivery



Conclusion

Are you ready to take your AI journey to the next level? Pure Storage offers simple and flexible AI-optimized infrastructure that can flexibly adapt to the demands of analytics and AI—launching enterprises, research entities, and scientific organizations into a new era, enabling faster innovation and a competitive edge.

Find out more about:

- [AI Solutions](#): Accelerate time to insight, seamlessly
- [AIRI//S](#): AI Ready Infrastructure from Pure Storage, powered by NVIDIA
- [FlashBlade//S™](#): Simple, scale-out, storage built for AI demands

purestorage.com

800.379.PURE

