

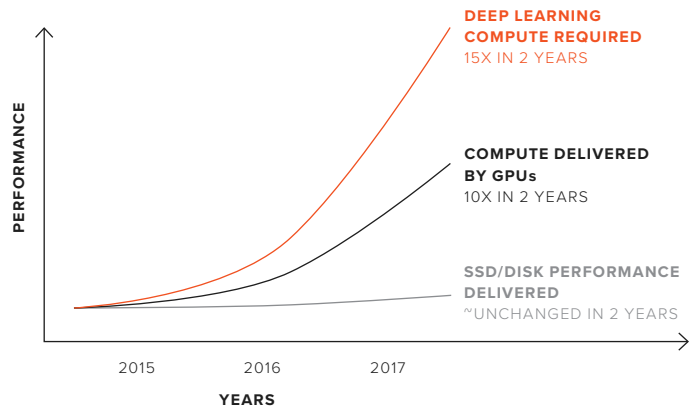
# MODERN MACHINE LEARNING WITH FLASHBLADE

## SUMMARY

Data is the fuel driving rapid innovation powered by artificial intelligence. Enterprises need modern data platform purpose-built for machine learning, accelerating insight while simplifying complex data pipelines in analytics.

### LEGACY, RETROFIT STORAGE NOT KEEPING UP

Artificial intelligence (AI) revolution today is powered by modern technologies, like deep learning algorithms and GPUs, that are massively parallel. However, legacy storage systems were largely built on decades-old building blocks, designed in the serial era, and the performance gap between compute and storage continues to grow.




### FLASHBLADE™ SCALE-OUT STORAGE REFLECTS THE SOUL OF AI

FlashBlade is the industry's first scale-out storage purpose-built for modern data analytics like machine learning. Fast, big and simple, FlashBlade is massively parallel at its core, offering unprecedented performance and simplicity for data scientists. Purity for FlashBlade operating software is built on a massively distributed key-value pair architecture. Like a GPU-accelerated server, FlashBlade is architected to accelerate parallel workloads, delivering the performance of 10 racks of disk in a 4U size.


#### NVIDIA DGX-1 SERVER

Delivers Performance of 10 Racks of CPU Servers



#### PURE FLASHBLADE

Delivers Performance of 10 Racks of Disk



**NOW YOU CAN:**

*Accelerate training by keeping GPU servers busy with data.  
Simplify workflow with modern, “tuned for everything” storage.  
Replace complex, fragile infrastructure of disks with Peta-scale capacity in 4U.*

**“TUNED FOR EVERYTHING” ARCHITECTURE**

Unstructured data has three qualities: volume is growing rapidly, it is random in its structure, and file/object size can vary from small to large. It’s no surprise that AI workloads reflect these qualities. Bigger training datasets means higher accuracy. Access patterns can be random. And files/objects can be small images and text to large streaming videos.


**FLASHBLADE IS ENGINEERED FOR UNSTRUCTURED DATA — BUILT TO DELIVER THE HIGHEST PERFORMANCE NO MATTER WHAT YOUR DATA NEEDS ARE**

- Delivers linear scaling **performance that grows with your data**, from TBs to PBs, to thousands of clients.
- Offers predictable, ultra-fast performance for **any access patterns, random or sequential**.
- Designed to deliver maximum performance, from **small, meta-data heavy to large streaming files**.

**POWERING ONE OF WORLD'S MOST POWERFUL AI SUPERCOMPUTERS**

A leading global, web-scale company has built a large AI supercomputer, powerful enough to be among the world's fastest systems. This impressive system is built with racks of NVIDIA DGX-1 servers and FlashBlade systems, designed to be well balanced for deep learning training. FlashBlade is engineered to deliver maximum performance in any deep learning training environment, reducing data processing time from days to hours.

**DGX-1**  
13K Images/Sec for each DGX-1  
Assume 115KB on average for images



**FLASHBLADE**  
Delivers 15GB/s to keep all 10 DGXs  
in rack busy with training data



**FLASHBLADE: DATA PLATFORM PURPOSE-BUILT FOR DATA SCIENTISTS**



**FAST**

Elastic Performance – up to 17 GB/s  
Always-fast, from Small to Large Files  
Massively Parallel from SW to Flash



**BIG**

10's of Thousands of Clients  
10's of Billions of Objects and Files  
1.5 Petabytes with Single IP, 4U



**SIMPLE**

Evergreen – No TB Re-Buys  
No Manual Tuning  
Just Add Blades for Performance



VISIT [PURESTORAGE.COM/ANALYTICS](https://www.purestorage.com/analytics) TO LEARN MORE