

# MAXIMIZE MACHINE LEARNING PERFORMANCE WITHIN A SINGLE, UNIFIED STORAGE SYSTEM



## Top Challenges

- **Data ingest requires massive, cost-effective storage combined with flexible access methods.**
- **GPU-accelerated computations need high throughput without latency.**
- **Accomodate varying workloads without relying on third-party solutions.**
- **Containerized pipeline stages require Docker and Kubernetes-aware persistent storage that scales quickly.**
- **Workload isolation is necessary when running multiple projects in parallel.**
- **End-to-end encryption required to protect personal and other sensitive data.**

## For TensorFlow Users:

- **Get up to a 30% boost in throughput by using our TensorFlow File System plugin.**
- **User-space driver reduces kernel mode transitions, and lowers CPU usage.**
- **Keeps costly GPUs fully utilized.**
- **Speed up the training and operational performance of your existing infrastructure.**
- **Do more learning iterations, faster.**

## Knowledge takes Performance (and Capacity)

The days of debating the benefits of Machine Learning (ML) are over. Now it's a race to see how quickly companies of all types can develop and deploy their ML solutions and gain a competitive advantage.

Unfortunately for many, the rush towards ML can highlight some glaring weaknesses in their existing storage infrastructure. Whether it's poor scalability, latency issues, or limited access methods, these problems can doom an ML project to failure. Further complicating matters is the fact that different stages of the ML pipeline place vastly different demands on the storage system.

Quobyte Software Storage removes these barriers by delivering the throughput, IOPS, scalability, and security to make that ML project a success.

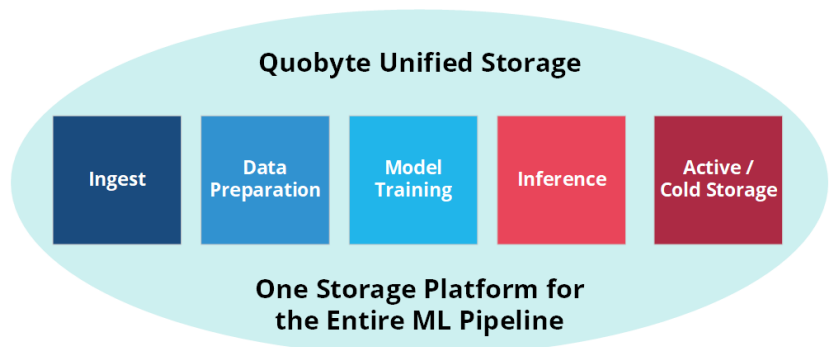
## Feeding the Pipeline

Because different stages of the Machine Learning workflow present distinct storage challenges—each with its own set of requirements—it's vital to use a storage system that meets those challenges. For example, it makes no sense to use a flash-only product to store ever-growing petabytes during the Ingest or Archive stages. But when the time comes to feed those data-hungry GPUs, the fastest networking and NVMe devices are the way to go. Quobyte never forces customers to compromise. The ability to tailor storage performance and space efficiency to match the workload is a hallmark of our solution.

## No Time for Point Solutions

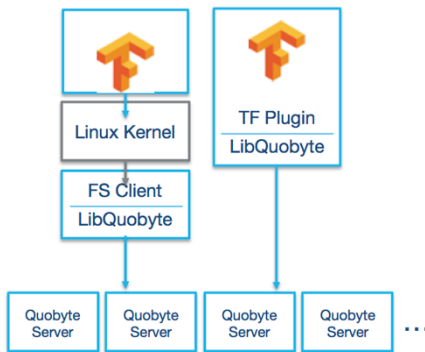
Few people would buy a car that only worked in the daytime and another for the night, so why would anyone want to cobble together two or three separate products to handle their ML storage requirements? But that is exactly what many vendors expect their customers to do.

Quobyte's unified storage eliminates the time-consuming and error-prone process of copying data from stage to stage. With our approach, data never has to leave the single file system throughout the entire ML lifecycle.



## Key Benefits

- Throughput up to 100Gbps per client, up to hundreds of gigabytes/s per cluster.
- Tens of thousands of IOPS with <200µs latency.
- Supports HDDs, SSDs, and NVMe devices in a single system. No need to offload long-term storage to external vendors.
- Erasure coding for space efficiency.
- Kubernetes CSI plugin and operator.
- Multi-protocol supports disparate data ingest sources.
- Independently scalable metadata, no bottlenecks.



## Get Quobyte Now!

Follow this link to our website

[www.quobyte.com/get-quobyte](http://www.quobyte.com/get-quobyte)

where you can download a 45 day eval copy of Quobyte and be up-and-running in an hour.

## Quobyte – The Data Center File System (DCFS)

Quobyte's next-generation file system unifies file, block and object storage for enterprise and scientific applications. Backed by a decade of research, its parallel file system core serves low-latency and high-throughput workloads within a single system.

## Enterprise Ready, Highly Configurable

Customers choose Quobyte's software storage solution for the operational ease and flexibility it offers. Features like user-definable data placement, policy-driven tiering rules, and hardware pinning give storage admins the freedom to optimize for their environment. Yet with all of this flexibility, the system is far easier to operate than a number of alternatives. And with its unlimited data capacity, non-disruptive live updates, and 1,000+ node scalability, you'll never outgrow your Quobyte investment.

## One Size doesn't fit All

In contrast to the outdated, HA pair-based approach to storage scaling, Quobyte software allows you to grow both capacity and performance linearly, letting you stay ahead of your increasing workload. And instead of forcing customers to run a secondary cluster network (and maybe yet another protocol), Quobyte's solution just needs a standard Ethernet infrastructure. Quobyte believes in freedom of choice, not vendor lock in. Customers are free to deploy on any hardware, and cluster nodes don't have to be identical. Furthermore, this freedom applies to the storage cluster OS as well, with support for all modern Linux distributions.

## Built for the Future

Yesterday's inflexible data center design is quickly giving way to tomorrow's highly responsive containerized world. With Quobyte, your storage will be ready. Not only does our software provide persistent container storage, it can also be run inside of containers. Combine that with Kubernetes support and a powerful RESTful API and you've got an ultra-responsive solution for rapidly fluctuating workloads. Quobyte also supports OpenStack's Cinder, Glance, Manila, and S3 interfaces, and is also compatible with Hadoop.

## Linear Scalability

Starting with four servers, Quobyte can be extended drive by drive and server by server. As a shared-nothing architecture, its IOPS and throughput scales linearly with every added resource.

## Lights-Out Operations

High availability and data protection is baked into the Quobyte software, because all hardware eventually fails. Together with extensive management functionality, Quobyte handles failures automatically and enables non-disruptive upgrades.

