



## Liquid Element LQD4500 Solution Brief

### Why Liquid

Unlock the full potential of your data with the Liquid Element LQD4500 IO Accelerator: the ultimate solution for high-performance, high-density storage that can meet even the most demanding workloads.

### Key Advantages

- » Unleash maximum productivity by powering through data-intensive workloads with lightning-fast performance.
- » Meet growing capacity demands efficiently without increasing footprint, with our high-density storage solutions.
- » Ensure data is safe and secure with advanced data protection and software RAID support.

# Do More Faster

## Unlock New Levels of Data Acceleration and Efficiency with Liquid Element LQD4500 IO Accelerators

### Overview

The challenge of enterprise computing is simple: data needs to be delivered to the computer before work can be done. As technology advances, CPUs are becoming more powerful and require even more capable storage for better results. Technologies like GPU computing increase this demand, as they are much more compute and data-intensive than traditional CPUs. The use of Private Clouds in the enterprise creates "compute islands" that make it increasingly difficult to get data to the workloads that need it, while minimizing the impact on other workloads.

To address these challenges, the storage industry has shifted from high-latency magnetic media to NAND-based media that provides high bandwidth, low latency, and minimal seek times. Additionally, new protocols and file systems have been developed to optimize access to these new, high-performance devices. In today's modern data center, PCIe-based NVMe storage is the standard technology for efficient data delivery across all modern compute workloads.

The Liquid Element LQD4500 IO Accelerator is a solution that delivers world-record data performance to satisfy today's most data-intensive workloads. The LQD4500 is a Gen 4.0 x16 PCIe NVMe device designed to maximize I/O with up to 4M IOPS, 26 GB/s throughput and only 10µs latency. It offers ultra-high performance density in a single PCIe slot, with up to 32TB on a single FHFL PCI device. Additionally, it has enterprise-class power fail protection, ensuring that data is persisted to flash in case of sudden power loss.

This solution brief will discuss the benefits of deploying the LQD4500 for modern compute workloads, specifically for AI/ML, Edge, Private Cloud, HPC, Media and Entertainment, System Integrators, and Web 2.0 and Cloud Providers.

## AI/ML

Artificial Intelligence (AI) and Machine Learning (ML) workloads vary, but they typically follow a standard compute and data workflow: deliver data to a compute host, perform calculations on it, and provide results. Challenges in these workloads are often found in the areas of data scale and latency, which have completely reshaped how to optimize this workflow.

- » **Performance and Scale:** ML and AI training require large amounts of data. As the size of the dataset increases, so does latency, which reduces the value of the results. The two common approaches to reducing data latency for AI/ML workloads are to leverage a high-speed data cache layer that copies data to the compute host or to use accelerator direct storage like NVIDIA Magnum IO™, where the compute accelerator accesses the data directly and bypasses the CPU entirely. With either method, both data locality and latency can limit overall workload efficiency. The LQD4500 is designed to serve in both functions, by either enabling a high-speed local cache or an IO accelerator leveraged directly by the GPU. With capacities of up to 30TB and internal bandwidth of 26GB/s, the challenges of keeping the data fed to the accelerator is a thing of the past.
- » **Cost Control:** Due to the unique performance demands of AI/ML, SATA-based storage can be expensive as more disks must be introduced to compensate for latency that increases at scale. The LQD4500 reduces the total cost of ownership (TCO) of AI-based solutions by delivering high performance density, reducing the need to add additional disks for performance, and saving valuable space in the data center.

## Edge Computing

Edge computing is creating new challenges for data management and analytics. With decentralized data sets and the need for near real-time analysis, technologies like vehicle automation, robotics, and traffic management are becoming more dependent on edge computing.

- » **Sensor Ingest:** As sensors become more advanced, they generate more data. To handle this increase in data, innovative solutions are needed. The LQD4500 is designed to provide high-performance I/O acceleration at the edge, allowing advanced sensors to process large amounts of data.
- » **Real-time Inference Analytics:** To optimize edge computing workflows, new requirements have emerged for processing data at the edge instead of relying on centralized processing. The LQD4500 delivers the performance needed for next-generation real-time inference and distributed learning.
- » **Performance Density:** Edge computing environments can often be cramped and challenging. With up to 32TB on a single PCIe card, the LQD4500 can operate at full performance in even the most demanding environments, allowing for large-scale data storage in a reduced footprint.

### Private Cloud

For private cloud use-cases, the primary metric for measuring value is the number of VMs that can fit on a single node. The more VMs that can be run on a single node, the more value it provides. Additionally, the more performance that can be delivered from a single VM means more types of workloads can be serviced with that private cloud.

- » **VM Density:** The Liquid LQD4500 provides leading class performance density with up to 32TB of capacity and 26GB/s of throughput. This allows more VMs to be run per node on I/O bound hypervisors, which improves the TCO of your solution.
- » **Private Cloud Scalability:** By using the LQD4500 as a pass-through device or via virtual filesystem, high-performance I/O acceleration can be provided to virtual workloads. This also allows for private cloud to service workloads that typically have I/O requirements that are not feasible with standard server designs, increasing the scalability of your private cloud.

### HPC Environments

Modern high-performance computing (HPC) environments provide a shared computing framework for research and product development across various disciplines. Researchers and administrators who use HPC are under constant pressure to process larger data sets in less time. The Liquid LQD4500 can play a critical role as a local data cache on a compute node or as high-speed shared file systems to optimize metadata or transient datasets.

- » **Accelerate Time to Insight:** HPC workloads are compute and data-intensive, so faster storage enables more capable HPC architectures. The LQD4500 provides high IOPS and throughput with low latency, making it the highest performance single card solution on the market for HPC workloads.
- » **Simple Reliability:** To implement RAID protection, typically 3-4 SSDs are needed. With the LQD4500, RAID can be implemented on a single PCIe card. This reduces the number of storage devices needed, saving costs and freeing up valuable PCIe slots for more storage or accelerators.

### Media & Entertainment

The Media and Entertainment workflow typically involves three phases: Capture, Process, and Deliver. Each phase is highly dependent on storage bandwidth. During the Capture phase, video streams are read-in and saved. The more video streams a single device can handle, the better the TCO. In the Process phase, the source file is read, processed, and a final version is written out. Faster read and write completion saves time and money. The Deliver phase varies by use case, but typically a streaming-type application is reading-in a file and sending it out to a client. More clients served per device means more revenue. By adding the Liquid LQD4500 as a high-speed data processor in the media pipeline, it can drive significant value for these types of workloads.

- » **Sequential Performance:** For streaming large amounts of data sequentially, higher sequential performance enables a greater number of streams supported by a single SSD device. With 4x the sequential performance of SSDs, the Liquid LQD4500 can support 4x the number of streams per device.
- » **Data Processing:** In the media and entertainment space, it is common for multiple designers to edit and post-process a scene at the same time. The LQD4500 provides a single card solution that can be installed in a designer's desktop, supercharging their day-to-day tasks and enabling them to design, render, and post-process at record speeds.

### System Integrators

System integrators are always seeking storage solutions that drive revenue and satisfy customers. Liquid offers a unique solution as the leader in IO acceleration, delivering more value to customers and integrators alike.

- » **Differentiated Offering:** System vendors can now offer unique solutions to their customers with higher performance. By offering the LQD4500 IO accelerator, system builders can provide their customers with the highest performing storage device on the market.
- » **Increased Profits:** With a competitive price compared to other solutions with similar performance, system vendors have the opportunity to increase their margins with the Liquid LQD4500. By providing differentiated systems with greater performance at a lower cost, they can offer better overall value to their customers.

### Web 2.0 and Cloud Providers

Today, service providers are under pressure to deliver more with less. Customers expect high performance, but are unwilling to pay more for it than for previous generation technology. Whether you're providing a DBaaS, mail service, web/streaming, or consolidated backups, the Liquid LQD4500 offers an optimal platform for creating top-notch services at competitive costs.

- » **DBaaS:** As a primary data store for relational or unstructured databases, the Liquid LQD4500 is an ideal solution for increasing overall TPS and driving optimal value from your design. Whether you're consolidating databases or creating a new top-tier solution, Liquid can help you exceed your goals.
- » **Mail Services:** To get the most value from a mail solution, you need the highest storage density possible in your environment. This translates to more users per mail server. The main challenge is typically I/O scalability. The LQD4500's ability to double or more the number of IOPS means your solution can easily handle mail tasks, allowing you to provide the performance your customers expect with the densities you need for optimal TCO.

- » **Web / Streaming:** Web content is largely driven by active, or database-created content. In these types of environments, every web server is also a database server or, at the very least, a database client. Adding the LQD4500 as a cache layer or as a direct data store can allow you to service more client requests with less hardware overhead. Additionally, since the LQD4500 offers the lowest latency storage on the market, your web server will reach new levels of performance and scale for your end users.
  
- » **Backups:** Service providers have long struggled with the challenge of providing backups to thousands of customers on a nightly basis. Typically, the approach is to add a few SSDs to the backup host to allow backups to land on a transient, high-speed location before moving it to less expensive media. Adding the LQD4500 to your backup hosts will allow you to drive higher backup densities, reduce overall backup times, and backup more servers with less hardware.

[Learn More at Liquid.com](https://www.liquid.com)