

# LA JOLLA INSTITUTE FOR IMMUNOLOGY OPTIMIZES COVID-19 RESEARCH WITH EXCELERO

*Excelero's NVMesh® software delivers 10X increased performance for exceptional ROI*


## CASE STUDY

### La Jolla Institute FOR IMMUNOLOGY

LJI is an independent research institution with a 30-year focus on analysis of immune system disease at the atomic, genetic, protein and cellular level – analyses that are fundamental across hundreds of diseases. Its team of over 200 Ph.D.-level scientists, informaticists, and research staff rely on high performance computing infrastructure to support timely analyses. As COVID-19 boosted research by its multi-lab Coronavirus Taskforce, and a tidal wave of data from its Krios cryo-electron microscope grew, LJI's storage demand was skyrocketing. With massive growth in its data sets, the IT team no longer could scale its infrastructure and deliver adequate performance with incremental additions of Just a Bunch of Flash (JBOF) storage resources.

With Excelero NVMesh as the centerpiece of a new storage architecture using elastic NVMe, LJI gained a 10x speed-up of storage performance on a mission-critical database, and an average 6.8x speedup across a variety of tasks. LJI now has a future-proof storage architecture that's hardware-agnostic, exceptionally performant, and lets nothing stand in the way of its life-saving research.

## Benefits of Elastic NVMe




- **Share NVMe resources across multiple GPU servers**
- **Access remote NVMe at local speed**
- **Exceed the performance and capacity limits of local flash on GPU servers**
- **Eliminate the need to copy data locally conserving time and drive endurance**
- **Datasets can be larger than what can fit inside the GPU Server**
- **Full CPU offload with Excelero's patented Remote Direct Drive Access (RDDA) technology**

## AVOIDING THE GPU BOTTLENECK

High-performance data storage is vital to medical research today, as investigators rely on data science in order to draw conclusions that may spark a breakthrough. To power them the IT teams at leading research organizations have embraced new GPU servers for many of their most computationally demanding workloads. Yet many research applications demand far more storage capacity than that available locally in GPU systems. Industry analysts have called this the “GPU storage bottleneck,” where the up to 2 million IOPS capability of the GPU vastly exceeds the 400K IOPS that local storage can feed it, and can increase time to insights.

At LJI, the IT team ran both a standard high performance computing (HPC) cluster and the GPU cluster, with storage comprised of JBOF deployments with conventional Flash drives. As demand for storage grew, LJI would add incremental GPUs and matched them with JBOFs storage resources, an affordable storage option which avoided often costly licensing requirements of traditional storage incumbents, and gave LJI the flexibility to deploy any future hardware it wished.

Expanded use of LJI's Krios high-end microscopy facility with its image file-intensive workload led to a rapid growth data volumes. By 2020 LJI's storage architecture had grown to 7 PB of raw data and nearly 900 million files. Running 7/24, the Krios facility alone generates 1-2TB of raw image data per day. Intensive sequencing workloads often involved 1 million reads per second, speeds that could swamp the limited local storage of the GPUs. LJI's IT team commissioned a system upgrade to enable both high storage IOPS and throughput to enable researchers to use Krios without IT constraints – while assuring the low latency that was essential to rapid response rates.



*“The storage system's critical feature is its custom database application, and Excelero has boosted transactions per second by 10X with that database alone. It's simply phenomenal.”*

*Michael Scarpelli, senior director of information technology, La Jolla Institute for Immunology*

## POWERING GPUS, AS CRYO-ELECTRON MICROSCOPY DATA SOARS

LJI's new storage infrastructure provided a harmonious answer to serving the superior processing speed of GPU servers in its GPU cluster as well as its HPC cluster.

The new high-performance storage system leverages Excelero NVMesh® to enable an accelerated software-defined storage stack that connects CPU and GPU compute servers to NVMe over the network and delivers performance that is unachievable in local storage deployments alone. While the Excelero solution is hardware agnostic, LJI utilized KIOXIA (formerly Toshiba Memory) data center NVMe® SSDs for system storage.


The architecture allows LJI to virtualize NVMe storage assets and deliver them as if they were local and with limitless scale. NVMesh is the elastic block storage engine that now speeds the revamped file server environment. This ensures the many complex interactions with petabyte-scale data sets will not be bottlenecked by storage infrastructure.

### Excelero NVMesh Storage Performance Testing at LJI

*"With the Excelero-driven system, we never even have to worry about performance. The system just flies."*

*Michael Scarpelli, senior director of information technology, La Jolla Institute for Immunology*

Task	Performance Increase
Synchronous Metadata Inserts	7.25x
Asynchronous Metadata Inserts	9.6x
Synchronous File Writes:	6.3x
Asynchronous File Writes:	4.1x
<b>Aggregate Performance</b>	<b>6.8x</b>



## ELASTIC NVMe PERFORMANCE, AT SCALE

Early testing showed that Excelero's NVMesh provides between 4.1x to 9.6x greater performance, depending on specific storage task, with this database application. On average NVMesh performs 6.8x faster than other options.

Based on workload, the Excelero storage architecture at LJL enables up to a 10X improvement in price/performance. Because NVMesh is hardware-agnostic, LJL now has a future-proof storage architecture enabling the rapid introduction of modern NVMe, network and server technologies without delays often introduced by rigid appliance-based solutions.

By incorporating Excelero NVMesh into its GPU cluster and storage architecture, LJL can now achieve the highest-performance at any scale, with the best price/performance. It has the flexibility to grow, and deploy whatever hardware it wants, knowing NVMesh will support it. Backed by its new infrastructure, LJL can fully expedite its mission by optimizing its data storage.

*"We were looking for blazing storage IOPS and throughput to enable researchers to use Krios without IT constraints, and we certainly found it with Excelero."*

*Michael Scarpelli, senior director of information technology, La Jolla Institute for Immunology*

### NVMesh features for GPU

- NVMesh unifies remote NVMe devices into a logical block pool that performs the same as local NVMe flash
- NVMesh allows full utilization of the IO/s and bandwidth capabilities of NVMe drives across a network
- Nvidia DGX-1 and DGX-2 can use their massive network connectivity to access remote NVMe logical volumes, with redundancy if desired!
- MUCH faster than local SATA SSDs
- Larger shared pools than possible within the platform
- Other GPU optimized systems can access remote NVMe at local latencies and bandwidth
- Random IO characteristics of NVMe preserved, achieving 10s of millions of potential IO/s at minimal latencies