



CASE STUDY

ICM Brain & Spine Institute Breaks the Storage Bottleneck for Life Sciences Computing

OpenFlex™ Brings Higher Resolutions, Faster Time-to-Discovery for Clinical Imaging



"Our scientists don't want to worry about storage technology, and with OpenFlex, they don't have to. Now they have fast, low-latency access to imaging data, in up to four times the resolution than before. The shared storage with NVMe-over-Fabric is there when they need it, so now our scientists can focus on using the data to advance their research, which has greatly increased their productivity and efficiency."

Caroline Vidal
CTO, ICM

Challenge

At the Institut du Cerveau et de la Moelle Épineière (ICM Brain & Spine Institute), scientists are developing new weapons in the fight against neurological disorders. To do it, they need to capture and analyze patient data from digital light sheet microscopes, MRI systems and other clinical imaging tools. Too often, however, investigators were hampered by an aging data storage infrastructure that could not keep pace with their needs.


For example, as ICM added new digital microscopes capable of capturing higher-resolution images, storage presented a constant bottleneck. The latest light sheet instruments can generate up to 2 terabytes of data every hour, with individual file sizes up to 400 gigabytes, far more than each workstation's local storage can accommodate. As a result, investigators were forced to work with lower-resolution files—reducing the information at their disposal to analyze disease processes. Even working with smaller file sizes, the workstations had to continually archive imaging data to secondary storage in the central data center. It could take four hours to copy data to storage—and the same time to retrieve it. This created significant delays at each step in the analysis pipeline, where scientists had to wait on the storage infrastructure to continue their work.

ICM's IT team considered adding local storage at microscopy workstations, but there wasn't enough physical space in the buildings. Additionally, researchers' storage and IT needs can change frequently as investigations wind down and new ones begin. It's common, for example, for different research teams (located in different buildings, using different types of equipment) to need bursts of several terabytes of additional storage for just a few days or weeks. Providing that kind of flexibility and elasticity was immensely challenging.

Solution

ICM IT leaders worked with their longtime technology partner, 2CRSi, to reimagine ICM's storage architecture. Together, they identified a new storage innovation that could provide the perfect combination of performance and flexibility: OpenFlex Composable Infrastructure from Western Digital.

OpenFlex brings state-of-the-art NVMe™-over-Fabric (NVMe-oF™) technology to the enterprise. It allows organizations to replace aging storage technologies with a high-performance networking fabric that takes full advantage of the latest, highest-performance flash storage media. With NVMe™-oF, ICM can provide I/O performance comparable to direct-attached flash arrays. At the same time, they can keep storage disaggregated from compute and allow applications to share a common pool of capacity regardless of location. ICM can use a single, centralized OpenFlex platform in the data center and share that capacity among dozens of workstations across the campus. Today, the system supports 10 light board workstations, but in the future, it will support 50 or more.



Western Digital's OpenFlex Composable Infrastructure Storage

To provide the foundation of the new storage architecture, 2CRSi recommended Western Digital's OpenFlex Composable Infrastructure, initially using five 15-terabyte¹ fabric-attached devices connected via RDMA over Converged Ethernet (RoCE v2). The platform supports up to 256 namespaces in each device, giving ICM the ability to provision over 1,000 distinct pools of storage to a large number of servers as needed. ICM can double the number of devices per platform over time as their requirements grow, scaling each platform up to 614 terabytes capacity within the same 3U form factor.

ICM uses OpenIO open-source object storage software to manage the infrastructure, providing erasure coding, multi-tenancy capabilities, data protection and other storage services. Using the OpenFlex Open Composable API, the IT team can easily automate the creation and management of composable storage using RESTful webservices—instead of having to manually contend with dozens of local storage servers.

Benefits

With OpenFlex, ICM investigators can now work with the same ultra-fast storage performance they would get from local NVMe-attached storage at each workstation, in a more efficient, flexible architecture. Effectively, ICM can now give researchers self-service storage provisioning and management on demand. Scientists can continually push their work forward, without having to worry about or wait for storage. And, using the Open Composable API, ICM's IT team was able to build these powerful new storage services quickly, with minimal effort.

For ICM's IT department, OpenFlex also means no longer having to try to squeeze in local storage close to workstations, or manually move storage around campus as projects change. Instead, they can distribute storage capacity over ICM's Ethernet network wherever and whenever it's needed, without sacrificing performance.

Western Digital.

5601 Great Oaks Parkway
San Jose, CA 95119, USA
US (Toll-Free): 877.816.5740
International: 408.717.6000

www.westerndigital.com

¹ One terabyte (1TB) is equal to one trillion bytes. Actual user capacity may be less due to operating environment.

©2020 Western Digital Corporation or its affiliates. All rights reserved. Produced 5/20. Western Digital, the Western Digital logo, and OpenFlex are registered trademarks or trademarks of Western Digital Corporation or its affiliates in the US and/or other countries. Please visit our website, <https://www.westerndigital.com/products/storage-platforms#composable-infrastructure> for additional information on product specifications.

OpenFlex provides:

- **Breakthrough performance:** Researchers can now access high-performance storage in whatever capacity they need, at the maximum throughput that ICM's 50-gigabit/second Ethernet links can deliver. The solution also assures the low latencies needed when working with high-resolution image files—ICM reported 34 microseconds or less for most storage operations.
- **Faster time to discovery:** OpenFlex performance has a direct impact on ICM investigators' lifesaving work, eliminating delays and interruptions in the analysis pipeline. ICM scientists can analyze and verify more clinical images per day, in 4x the resolution than was possible before, and quickly retrieve archived datasets when needed. As a result, they can advance their understanding of neurological disease processes and develop novel therapies more quickly.
- **Architectural flexibility:** With OpenFlex, ICM's IT team can quickly and easily allocate storage to meet any investigator's need and resize and reallocate storage volumes on demand. This flexibility will be essential as the institute adds more microscopes and other instruments in the coming years, continually driving up the resolutions and volumes of clinical imaging data.
- **Fast, simple management:** Due to the complexity of ICM's previous storage infrastructure, any changes required the IT team to manually check and recheck all systems and access methods. OpenFlex provides allows for much faster, simpler maintenance. The team can manage the entire solution using Puppet or simple command line instructions, and easily change or duplicate configurations. Even if ICM adds multiple OpenFlex instances in the future, it will be easy to configure and use the systems to share storage instantly across the campus.
- **Lower costs:** With storage centralized in the data center, ICM's IT team can manage the solution much more easily. Maintenance tasks that used to take hours can now be handled in minutes—a significant operational savings compared to managing dozens of storage servers distributed across the campus. Additionally, because storage is centralized, any new microscopes added can use the same pool of capacity, eliminating the ongoing capital expense of deploying more local storage for each new device.

"We appreciate 2CRSi's collaboration on this project", says Caroline Vidal, "as usual, their technical team worked proactively to select the appropriate and cost effective solution, while providing an advanced solution that combines expertise and synergies between user needs, IT constrains and suppliers capabilities."