

QF2 for Research Imaging

Qumulo File Fabric (QF2) is a modern, highly scalable file storage system that runs in the data center and in the public cloud. QF2 handles large and small files with equal efficiency and can scale to billions of files. Its performance, cost, real-time visibility and simplicity make it ideal for institutions engaged in cutting-edge research imaging.



Performance

Managing throughput and latency are critical for research imaging. Researchers want to view images in real time and are sensitive to delays. Viewing images, especially 3D images, can require many random reads from large files.

QF2's hybrid architecture with adaptive read caching provides the performance needed for random reads within large files, such as OME TIFF files. Read caching uses an access-based heat map to intelligently move data to the flash tier when needed. QF2 provides high performance even with 3D image files that require random small block reads to access the image the researcher wants to see.



Efficiency

QF2 uses 100% of provisioned capacity for file storage, in contrast to the 70% or 80% of usable capacity that legacy storage systems recommend. QF2's efficient use of disk space decreases the data footprint and saves not just on the cost of disks but on infrastructure costs, such as power and cooling.



Real-time visibility and control

QF2's real-time visibility and control are extremely useful for managing research imaging workflows. Administrators can assign real-time quotas, so they're always in control of how resources are allocated among different projects. Rogue processes that can use up a storage system's resources are easy to identify and stop.

The capacity explorer and capacity trends tools give up-to-the-minute information about how storage is being used now and how storage has been used over different periods of time. With these tools, administrators can give researchers realistic numbers when they include storage costs in their grant applications.



Simplicity

Research institutes usually have a small IT staff. The complexity of setting up legacy storage systems and the lack of insight they offer make them expensive systems in terms of management. From the moment QF2 is unboxed to when it can start serving data is a matter of hours, not days. (QF2 for AWS has almost instantaneous setup.)

Universal-scale file storage

Scales to billions of files

Use any mix of large and small files and store as many files as you need. There is no practical limit with QF2's advanced file-system technology.

Your data anywhere

QF2 runs in the data center and in the public cloud. Continuous replication moves the data where it's needed when it's needed.

Highest rated support

Get help fast from our team of storage experts with your own Slack channel.

Cloud-based monitoring

QF2 proactively detects potential problems, such as disk failures. You can also access historical trend data about how your system is being used.

Complete REST API

Use the QF2 REST API to build and manage a modern application stack. It's the future of infrastructure, available today.

"We used to say 'just add more capacity' but storage is so much more complicated and mission critical than it was 10 years ago that, for us to effectively grow, we need to treat storage as a finite and precious resource and manage it accordingly. We can do that with QF2."

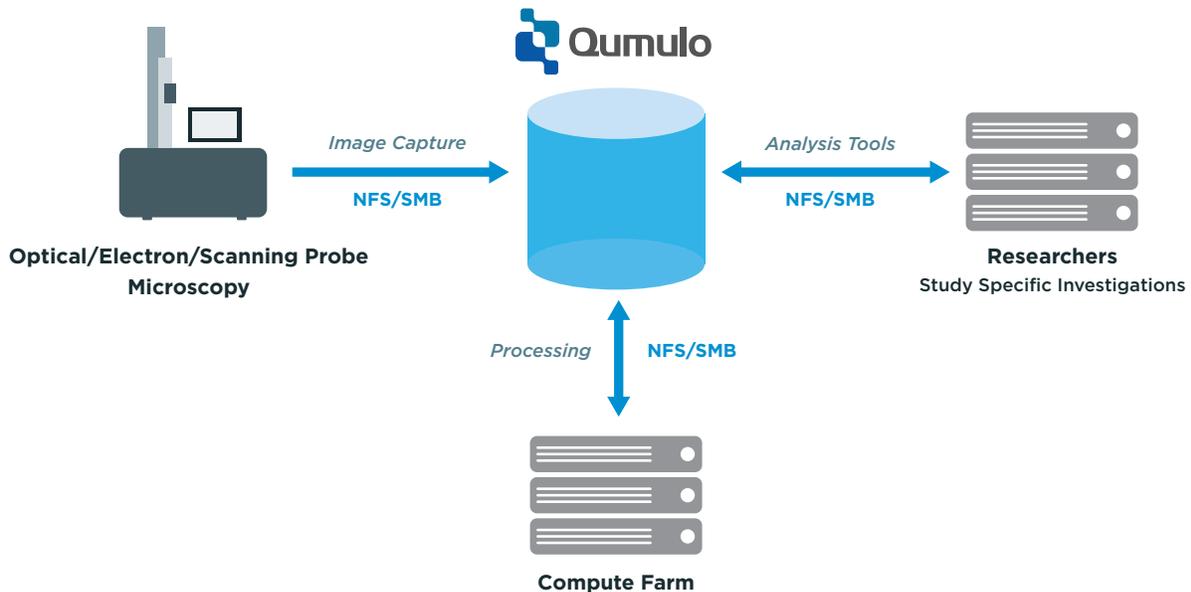
Nick Rathke

Associate Director
Information Technology
Scientific Computing and Imaging Institute,
University of Utah



Research Imaging Workflow

Here is an example of a research imaging workflow.



The images generated by the microscope are transferred to the QF2 cluster. Typical data sets use the standard proposed by the Open Microscopy Environment (OME), the OME-TIFF format. An OME-TIFF data set stores the 2D image planes that make up a 3D composite image within one multi-page TIFF file, or across multiple TIFF files. A metadata block describing the dataset is embedded in each TIFF file's header.

The representation of 3D as 2D image planes impacts storage access patterns: random reads from within an OME-TIFF file are common. Each component of the microscopy pipeline gets its own TIFF file. Within an OME-TIFF file there is an OME-XML metadata block which provides all the compatibility of TIFF while allowing you to take advantage of the rich metadata defined by OME-XML and the OME Data Model. QF2's read-caching algorithm means that frequently accessed files stay on the SSD for fast access.

Processing occurs on the compute farm and, at the same time, researchers can carry out their analyses on their own workstations. QF2 performance ensures concurrent access of the data by researchers and compute resources.

About Qumulo

Qumulo is the leader in universal-scale file storage. Qumulo File Fabric (QF2) gives data-intensive businesses the freedom to store, manage and access file-based data in the data center and on the cloud, at petabyte and global scale. Founded in 2012 by the inventors of scale-out NAS, Qumulo serves the modern file storage and management needs of Global 2000 customers. For more information, visit <http://qumulo.com>.

