


# QF2 for Data Analytics

Qumulo File Fabric (QF2) has the performance, scalability and enterprise features required for data analytics workloads. QF2 provides a single namespace that eliminates data silos and makes multiple copies of the data unnecessary.



## Get results faster

QF2 has better sustained read throughput than direct-attached storage for analytic workloads. Operating over today's fast networks, QF2 outperforms HDFS infrastructure. The performance edge of QF2 comes from its hybrid SSD/HDD architecture and its advanced, distributed file-system technology.



## Buy only the storage you need

QF2 allows you to scale storage and compute independently. With QF2, you avoid overprovisioning because you only buy the storage you need, regardless of how your compute cluster grows.

QF2's efficient data protection, based on erasure coding, gives you more usable capacity on your storage system. You'll save money on disks as well as infrastructure costs such as power and cooling.



## Real-time visibility and control

QF2's real-time visibility and control provides information about what's happening in the storage system, down to the file level, no matter how many files are in the system. System administrators can apply quotas in real time. The capacity explorer and capacity trends tools give IT the information it needs to plan sensibly for the future and not waste money because of overprovisioning.



## Cloud and on-prem

Organizations that want to move some of their data analytics to the cloud can take advantage of QF2 for AWS. QF2 has the highest performance of any cloud offering and is the only file storage system in the cloud with a full set of enterprise features, such as multi-protocol support and real-time visibility.

QF2 uses continuous replication to move data where it's needed, when it's needed. Continuous replication means you can easily transfer data from your on-prem QF2 cluster to your QF2 cluster in AWS, perform your computations, and then transfer the results back to the on-premises storage.

## 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.

### Single namespace

QF2 provides a single namespace for all the data, which eliminates copying and data silos.

### 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.

### Out-of-the-box simplicity

Skip the pain. From the moment QF2 is unboxed to when it can start serving data is a matter of hours, not days. QF2 for AWS can be set up instantaneously.

### Complete REST API

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

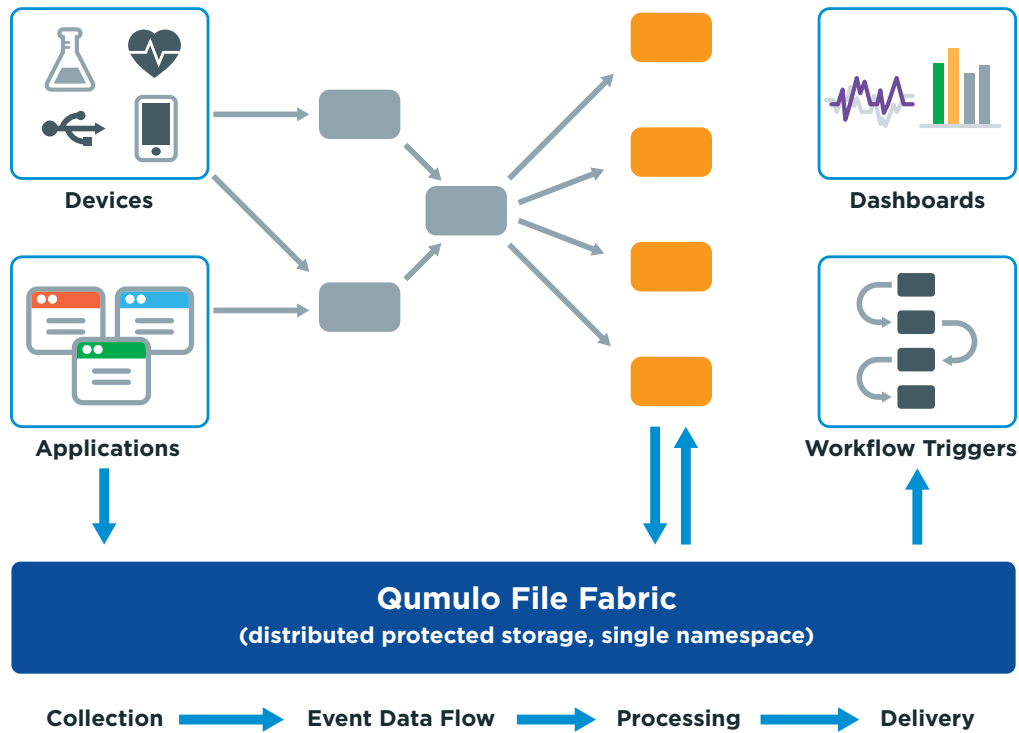
*"Managing data with QF2 is so simple it's hard to describe the impact. It has given us tremendous ROI in terms of time saved and problems eliminated, and having that reliable storage we can finally trust makes us eager to use it more broadly throughout the company."*

**John Beck**  
IT Manager  
Hyundai MOBIS



## Data analytics workflow

Here is an example of a streaming data analytics workflow that shows QF2 as the central, storage for the entire process, from ingesting the data to displaying it and acting on it.



In this example, some input comes from devices, such as cell phones, scientific instruments, autonomous vehicles and serial devices. It can also come from applications, which typically store their data in QF2 and then send a link to the event data flow software packages. The compute resources process the data and both store and retrieve files from QF2. Finally, the results are delivered and either displayed as information on a dashboard or used to trigger a particular action, such as a security alert.

QF2 provides a single repository for all the data. A centralized repository means that storage can be scaled independently of the compute resources. You only buy the storage you need, regardless of how the compute cluster grows.

A single file-storage repository eliminates data silos and the need for multiple copies of the data. It simplifies both the architecture and the workflows, and provides a cost-effective way to store and manage valuable information.

## 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>.

