CURE LIFE SCIENCES RESEARCH WOES WITH SCALABLE STORAGE SIMPLICITY

Q

Quobyte

Top Challenges

- **Proliferation of faster, less expensive Next-Generation Sequencing systems**
- **Increased sequencing runs to maximize Phred scores**
- **Far more people being sequenced**
- **Automated microscopes producing huge image files**
- **Long-term retention to support ongoing shared research**
- **Tightening regulatory compliance and security requirements**

> *We are on the threshold of a massive wave of genomic utility — a transition from the more academic research focus of today to a transformed medical landscape where sequencing is a routine part of patient care....*
>
> *...Combine all this with a stronger toolset for genome editing (think CRISPR), and we will see a huge leap in sequencing demand, with a market easily 1,000 times bigger than what we see today. That demand will drive the creation of cheaper, better, and faster solutions."*
>
> Andy Hollinger, Broad Institute
> September 13, 2016 blog post

Genomic and microscopy breakthroughs

Ever since Frederick Sanger developed DNA sequencing methods in the 1970s, researchers have sought to speed up the process. When Next-Generation Sequencing (NGS) arrived in the early 2000s it redefined sequencer performance, dramatically driving down the cost per base pair.

Similarly, modern microscopes have benefited greatly from technological advances. Evolved from the optical, manually operated scopes of the past, today's researchers feed automated imaging microscopes via Automated Tissue Processors (ATP) that slice and mount specimens.

This progression brought with it a tremendous decrease in the unit cost per test. But those millions of FASTQ files need high-IOPS during sequence alignment (i.e. BLAST, Bowtie, MAQ) runs. And high-res microscopy (i.e. epifluorescence, confocal) images have become truly massive. Consequently, storage admins in Bio IT are facing an explosion in the performance and capacity requirements to support modern research centers.

Capture it, process it, save it, share it

Capable of supporting dozens or even hundreds of sequencers simultaneously, a Quobyte storage cluster can easily handle billions of base pairs per hour. As processing of these files moves to the IO-intensive alignment phase, our parallel file access can shave time from this long-running task. Should you need to run some big data analytics, Quobyte storage offers a driver for Hadoop.

Much of today's genomic research is freely shared via facilities such as GenBank, SRA, GEO, ENA, and others. As these repositories continue to grow, so too does the need for economical, long-term storage. Quobyte's policy-based automation lets you move files to space-efficient, erasure coded storage based on the criteria you select. And sharing this valuable data over the web is simple with our stateless S3 proxy. Need to keep proprietary research safe? Multi-tenant support and device-level control gives storage admins total flexibility over where that data resides.

One size doesn't fit all

In contrast to the outdated, HA pair-based approach to storage scaling, Quobyte software allows you to grow both capacity and performance linearly, letting you stay ahead of your increasing workload.

Quobyte believes in freedom of choice, not vendor lock in. Customers are free to deploy on any hardware, and cluster nodes don't have to be identical. This lets IT perform hardware refreshes on their own schedule.

Key Benefits

- POSIX file system for seamless integration
- High IOPS, consistent sub-millisecond latency
- Parallel I/O delivers unrivaled throughput
- "Lights-out" data center resiliency; self-healing
- Near-perfect linear scaling
- 100% hardware and kernel independent
- Hadoop, Docker and OpenStack support
- Erasure coding option, perfect for analytics and long-term workloads
- End-to-end CRC assures data integrity, eliminates "bit rot"
- Policy driven data placement and tiering
- Integrates with Kubernetes, Rancher, and Mesosphere

Get Quobyte Now!

Grab our Evaluation Guide, download Quobyte and be up-and-running in one hour.

Click here »

Keeping data access fast and reliable

Delivering hyper-performance storage while maintaining data availability despite disk or network failures takes a holistic design approach. Customers can trust Quobyte storage because it was built to handle real-world failures and still remain operational. True end-to-end checksums of every data block assures data integrity. Ultra-low latency and fast meta-data operations keep up with the toughest small-file workloads. Hardware monitoring will offline an errant disk before it causes corruption, and self-healing heuristics route around node failures.

Enterprise ready, highly configurable

Customers choose Quobyte's software storage solution for the operational ease and flexibility it offers. Features like user-definable data placement, policy-driven tiering rules, and hardware pinning give storage admins the freedom to optimize for their environment. Yet with all of this flexibility, the system is far easier to operate than a number of alternatives. And with its unlimited data capacity, non-disruptive live updates, and 1,000+ node scalability, you'll never outgrow your Quobyte investment.
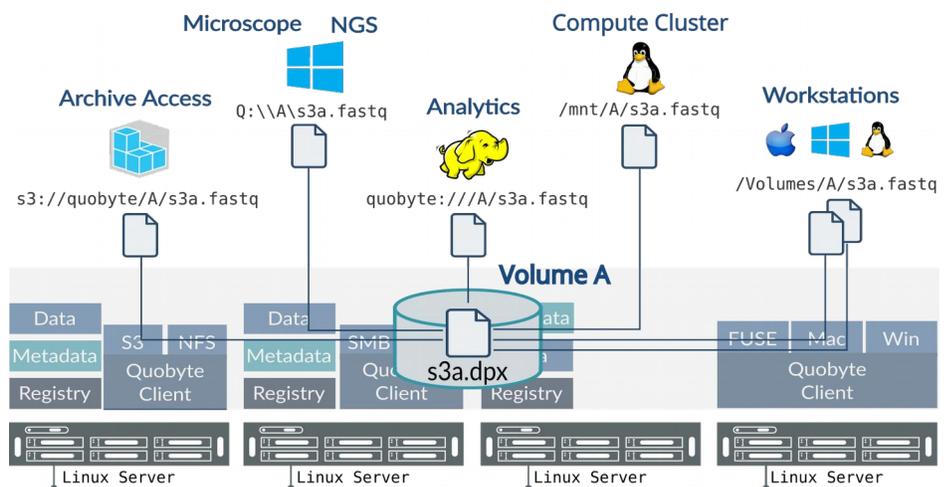
Built for the future

Yesterday's inflexible data center design is quickly giving way to tomorrow's highly responsive containerized world. With Quobyte, your storage will be ready. Not only does our software provide persistent container storage, it can also be run inside of containers. Combine that with Kubernetes support and a powerful RESTful API and you've got an ultra-responsive solution for rapidly fluctuating workloads. Quobyte also supports OpenStack's Cinder, Glance, Manila, and S3 interfaces, and is also compatible with Hadoop.

Quobyte®
Data Center File System.™ Fast and Reliable Software Storage.