

Quobyte Provides Down-to-Earth HPC Storage Solution for 3vGeomatics



Summary

Needs/Challenges: 3v Geomatics (3vG) is an innovative provider of Interferometric Synthetic Aperture Radar (InSAR) technologies that utilizes radar satellite images to detect and measure ground and infrastructure displacement across large areas. The company's exponential growth over the past 5 years has led to significant challenges in how best to store and access massive amounts of data for HPC processing and analysis.

Solution: Quobyte® Data Center File System

Platform: HPC Data Center and Grid Engine

Use Case: Commercial HPC, Oil and Gas, Mining

Key Benefits

- High through-put, low latency scalable performance
- Elimination of NFS bottlenecks
- Elimination of storage silos via many file access protocols
- Reduced management complexity for maximum operational efficiency
- Nearly instantaneous response to changing storage requirements
- Storage simplicity at scale

“The throughput is much better than we ever had. If it's a read or write-heavy program, we're seeing an order of magnitude improvement.”

Joe Chapman

IT Manager at 3vGeomatics

Background

Founded in 2007, 3v Geomatics (3vG) is an innovative provider of Interferometric Synthetic Aperture Radar (InSAR) technologies for a broad base of users in industries such as mining, oil & gas, urban infrastructure, transportation, water management and others. By listening to its customers' requirements and investing heavily in R&D, 3vG has improved the automation and scalability of InSAR to be able to deliver displacement data that could be better integrated into its clients' existing workflows.

3vG utilizes radar satellite images to detect and measure ground and infrastructure displacement across large areas. Each image contains more than 100 million pixels, and 3vG's processing chain measures changes at each pixel across dozens or hundreds of images acquired over time. By providing near real-time information based on satellite images obtained just hours earlier, 3vG delivers precise measurements of where, when and how much displacement has occurred – measuring millimeters to centimeters of displacement from a distance of 600 km – helping clients proactively identify and monitor geohazards to improve safety and prevent large-scale environmental failures.

The Challenge: Improving 3vG's processing times on large-scale radar images was required to help prevent environmental accidents, improve safety, and maintain operational success

The company's exponential growth over the past 5 years has led to significant challenges in how best to store and access massive amounts of data for HPC processing and analysis.

"The challenge is really that the amount of data you get from a single satellite image can be on the order of several gigabytes," said Dave Mackenzie, Operator "We work with 50 plus sites at the moment and growing every day. We'll build a stack of images for any given site, so we might have tens to hundreds of images for every site, each of which are on the order of several gigabytes when they're raw. Once we receive that raw compressed data, we will expand that out to be up to on the order of 10 gigabytes per image within that stack. So the challenges of dealing with that quantity of data put a pretty heavy load on whatever storage system we choose to use."

In 2018, 3vG hired Joe Chapman as its systems engineer to find a storage system that could better handle its unique requirements. The company had a 1.2 PB ZFS server on Linux (ZoL) as its main storage system that distributed its data through NFS mounts. With an HPC data center in-house and a grid engine to manage it, the company was forced into implementing arbitrary I/O load limits for all of its nodes and processes to overcome the limitations on its slow NFS or the company would experience constant lockups.

"3vG had been struggling with storage since the dawn of time for this company," said Chapman. "They'd tried several different storage systems including Lustre, Gluster, and mhdfs before settling on ZoL by the time I got here. We're an InSAR company that now has to be HPC experts out of necessity."

After considering a number of options that were ultimately determined to either be too slow, not designed for performance, or prohibitively expensive, 3vG decided to implement the Quobyte Data Center File System.

"We required fast access given that we have a significant amount of raw data," said Chapman, "and via our processing chain, the size of our intermediate data files balloons before we get to final products. I looked at many different storage systems – some of them were open-source, some were proprietary software, some of them were all-in one hardware and software appliances. When I found Quobyte's HPC product, I thought, 'This seems to do everything we need.'"

The Solution: Quobyte's efficient software solution helps overcome many HPC storage challenges

Quobyte is a high-performance, parallel file system that delivers the performance necessary – whether it be for workloads that require high throughput, parallel processing, small file operations in OpenFOAM, or large sequential file operations. It also scales, just add the resources needed, Quobyte takes care of the rest – and saves valuable admin time. Quobyte facilitates the entire HPC workflow – no more silos, no more tedious and complex capacity planning, and better economics thanks to operational efficiency at scale.

Quobyte software provides non-stop, high-performance storage for the most-demanding workloads using economical, commodity hardware. Engineered to be a complete solution, the product does away with operational complexity, replacing it with scalable operations that require fewer human resources. Linear scaling, run-time configuration flexibility, and real-time performance monitoring means responding to changing storage requirements is nearly instantaneous. Support for rolling upgrades and non-disruptive cluster expansion removes the burden of planned outages. Quobyte storage offers the freedom to choose how data is stored, and delivers that data to clients over the broadest set of access protocols, all while maintaining consistent access control. Customers are able to combine the benefits of flash with the storage efficiency of hard drives, the product never forces the use one when the other will do. Quobyte delivers storage independence.

The Results: Faster processing times, easier management helps 3vG deliver InSAR results to customers without delay

After more than a decade in business, 3vG continues to have an entrepreneurial mindset, a customer-centric focus and a driving desire to deliver technology that has a positive impact on the world. As it continues to introduce InSAR into new sectors and new geographies, the company pushes itself to constantly propel the boundaries of InSAR's capabilities. This leads to innovative approaches and customized products.

One such product that was recently placed into production for its pipeline customers is currently utilizing Quobyte for processing images with throughput rates significantly better than its previous storage systems.

"The throughput is much better than we ever had on ZFS," said Chapman. "It's night and day. If it's a read or write-heavy program, we're seeing an order of magnitude improvement over ZFS and NFS and, more importantly, no more NFS lockups!"

Chapman was also impressed by the ease-of-management of Quobyte. After filling some empty bays with drives, he was able to do a quick initialization by logging into the Quobyte interface, clicking create and the drives were readily available.

"It was so stupid easy. I've had to do that for our ZFS server and because it's only one server, if I need to open it up for any reason, it becomes an after-hours thing. We need to do a full office shutdown. The farm is off, now I'm working late. 'Oh my God, I hate this!' On the management side, Quobyte is so much better. I really like Quobyte's unified approach to storage. Clients can access files via a myriad of protocols such as the native client through FUSE, NFS, SMB, or even S3. There is no need for a third party object store archive like some other solutions require. We've yet to even unlock the full potential with a hot flash tier."

3vG currently has twelve servers running Quobyte software with another four being prepared at a second location for disaster recovery purposes.

"Our customers use the data we give them to make decisions on the ground day-to-day," said Mackenzie. "As a high data throughput company, we've really developed a fast, 24-hour turnaround on reports after receiving a satellite image. We have deliverables due pretty much every day and we can't miss one. So a solution like Quobyte that does all the processing and everything 10 times faster than what we saw before, it's awesome."



Quobyte[®]
Data Center File System.™

Quobyte Inc.
4633 Old Ironsides Drive
Santa Clara, CA 90504
info@quobyte.com
T: 650-564-3111