Qumulo

Falcon's CREATIVE GROUP

Creating Immersive Experiences at Falcon's Creative Group

Falcon's Creative Group, in Orlando, specializes in immersive experiences for theme parks, museums, hotels, zoos, and more. It was founded in 2000 as Falcon's Treehouse by Cecil D. Magpuri, who was formerly a Creative Director for Universal Creative in Orlando.

The company has three divisions, Falcon's Treehouse, Falcon's Licensing and Falcon's Digital Media. The Digital Media division produces content that uses the latest advances in live action filmmaking, computer generated imagery (CGI), visual effects (VFX), virtual reality (VR), augmented reality (AR), mixed reality (MR), audio design, interactives, and gaming.

QF2 at Falcon's Creative Group

Saham Ali, Directory of Technology at Falcon's Creative Group explains the role Qumulo File Fabric (QF2) plays in Falcon's projects.

Saham joined Falcon's Creative Group in 2016 at a time when the studio was moving to more elaborate projects. He says, "The projects were getting bigger, resolutions were getting bigger, and file outputs were getting bigger." In particular, the company was well known for VR, even before headsets were mainstream. Falcon's was one of the first companies to render a 360-degree dome where objects could come out of screen space into audience space and back into screen space and could be viewed from any angle without breaking the stereo effect.

The projects were getting bigger, resolutions were getting bigger, file outputs were getting bigger.

To deal with the amount of time renders were taking, the company switched from CPU-based rendering to GPU-based rendering, but Saham found that the existing storage couldn't handle the workloads and was maxing out on IOPS.

Saham wanted a new storage solution, and his main criteria were scalability and simplicity. He says, "It couldn't be direct-attached storage, where we would be stuck with, say, 50TB and that was it. We also didn't want to dedicate an entire rack to get the throughput and IOPS we needed."

Saham first learned about QF2, a modern, highly scalable file storage system, while he was IT Director at The DAVE School (Digital Animation and Visual Effects). Afterwards, as a consultant, he started seeing QF2 at different studios. Now, with a chance to revamp the company's storage infrastructure, QF2 was what came to mind.

He and his team performed tests, evaluated the numbers, and proved to themselves that QF2 could give them the storage and performance they needed at a reasonable price.

He and his team performed tests, evaluated the numbers, and proved to themselves that QF2 could give them the storage and performance they needed at a reasonable price. They purchased four QC104 nodes connected with a 40 GbE switch.

When the system is under heavy use, Saham sees 2GB/s to 3GB/s reads and just under that for writes. When reads and writes are happening simultaneously, he sees both at around 2GB/s to 2.5GB/s.

QF2 came through on simplicity as well as performance. He says, "We built out the rack, and the Qumulo team came out to drop in the drives and turn everything on. We allocated the IPs, they configured the switch, and that was it. The biggest hiccup was on our end. We had a problem getting the core switch to talk to the 40 GbE switch, and that was an issue with the Dell switch CLI. We were up in a couple hours and writing files within the day. Now, it's so easy to add a node or allocate virtual IPs or get the analytics without having to do much. It's just simple."

We were up in a couple hours and writing files within the day. Now, it's so easy to add a node or allocate virtual IPs or get the analytics without having to do much. It's just simple.

As a bonus, Saham appreciated QF2's hybrid design. He says, "Qumulo knows that flash is the future. You understand the whole concept of RAM up front, then flash and then spindles. It was a huge selling point that you understood it and had a good working design."

The Church of the Holy Sepulchre

One of Falcon's most technically ambitious projects was Tomb of Christ: The Church of the Holy Sepulchre Experience for National Geographic. National Geographic documented the restoration of the church with millimeter-accurate LIDAR (Light Detection and Ranging) scans. They wanted to transform those scans into a 3-D experience they would show at their museum in Washington, D.C.

The massive amounts of data they had scanned were handed over to Falcon's to be replicated as 3-D objects. The scans themselves were billions of polygons with 8-10, 16K Textures. An average EXR file that had all the AOVs and was rendered in Maya would be anywhere from 2GB to 2.5GB per frame (the frame rate was 24 fps).

A compositing artist could use anywhere from two to six read nodes, each of them reading an EXR file. Saham says, "It was slower, and by increasing our throughput, we increased our productivity and ability to review iterations at a faster pace."

"As soon as we migrated to QF2, we never had to worry about how many nodes were accessing the storage at any given time and we didn't have to schedule when the artists could work. We threw everything we had at QF2 and nobody complained that things were slowing down."

We threw everything we had at QF2 and nobody complained that things were slowing down.

The end product is 10 minutes long. Audiences stand in front of a 270-degree screen that goes below the level of their feet and up to the ceiling. Visitors feel as if they're flying through the building. Although Falcon's had been doing dome renders for a long time, this was the largest- scale dome render that they'd been able to achieve in the relatively short timeframe of six months.

The Storage Network

Here is a simplified diagram of the storage network.

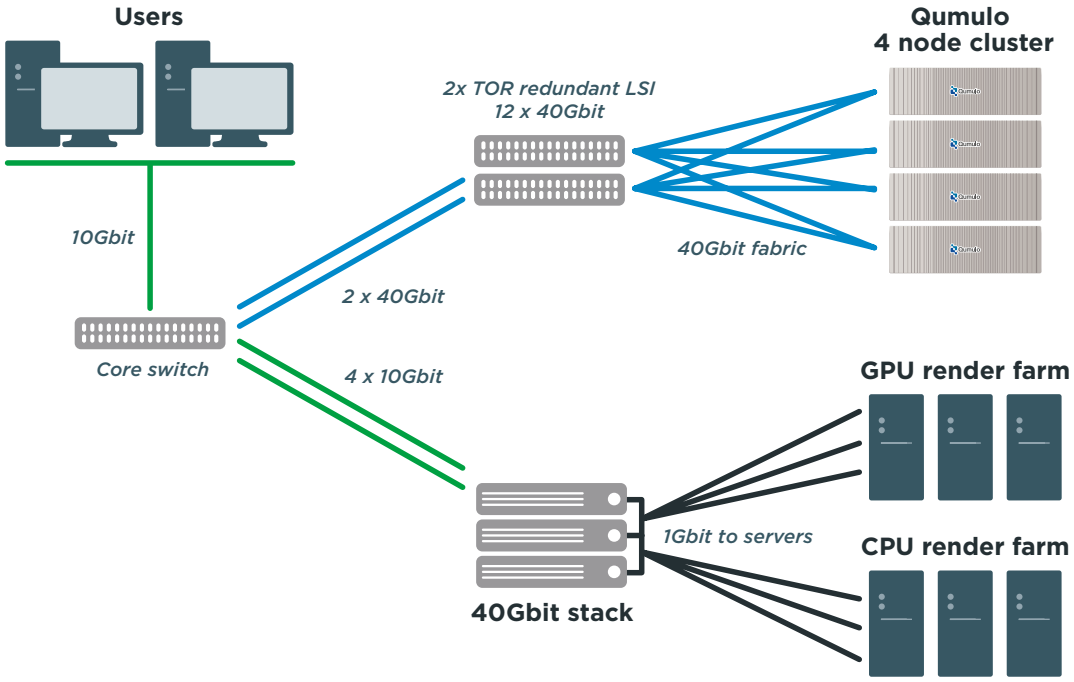

Everything is rendered with Maya using Redshift, a GPU renderer. From there, it goes to the compositing team.

The polygon problem

When Falcon's agreed to take on the National Geographic project, they knew there would be technical challenges. One was the complexity of the LIDAR scans, where a single section of the tomb could be made up of from 5 to 10 billion polygons. However, that was just the beginning. Because audiences would experience a seamless flythrough, multiple rooms would need to be loaded simultaneously, depending on what the camera would see.

Billions and billions of polygons would require many terabytes of RAM, so a brute force approach was impossible. The team needed to develop a caching mechanism that would drastically reduce the amount of RAM required.

Using our old storage system, it would have been difficult to meet the demand of deadlines, and QF2 kept us on schedule.

Saham says, "The tools we developed generated these very large caches, and they all sat on QF2. Using our old storage system, it would have been difficult to meet the demand of deadlines, and QF2 kept us on schedule.

Using QF2 Analytics

QF2 analytics, which are an integral part of the QF2 codebase, give Saham the insight he needs into his storage system. He says, "The real-time analytics are awesome because they let us know what's happening without digging around and figuring it out using third-party or handwritten tools.

The real-time analytics are awesome because they let us know what's happening without digging around and figuring it out using third-party or handwritten tools.

"If we want to know what's going on with the QF2 cluster, we only have to look at the integrated dashboard. We don't have to do a directory crawl and we get a real-time understanding of where all our storage space and IOPS are going. That type of visibility is huge.

"We can see which of our clients are the hottest and why because we have a granular understanding of what's happening. Is it the textures and the Maya render that's causing all the IOPS? Is it the cache files or the particle files? Finding out that information is simple."

The Future

Saham sees many opportunities for Falcon's and QF2. He's been thinking a lot about the cloud. He says, "We're definitely interested. The fact that QF2 runs in the public cloud was a selling point. The limiting factor right now for rendering in the cloud is cost because we are a GPU- centric house.

We're also thinking about using the cloud to extend our file system so that we can archive projects. Those are conversations we need to have. Whatever we do, QF2 will help us to continue to grow as a company."

Whatever we do, QF2 will help us to continue to grow as a company.