

Major Animation Studio Speeds to Finish Line on Latest Feature Film with IBRIX



MEDIA & ENTERTAINMENT

CASE STUDY

Digital effects houses and film studios have leveraged the computational power of high-performance clusters for rendering and other complex jobs. Over time, the size and makeup of the clusters have changed. Commodity-based clusters are replacing more costly symmetric multiprocessing 'big iron' machines, and node count has increased from the handful it might have taken a few years ago to thousands of processors today.

This case study discusses the challenges a California animation Studio had in making the summer 2006 release date of its latest film. Having produced six animated feature-length films over the past nine years, the Studio was a proven leader, but suddenly was left scrambling when faced with mounting production delays because frames were taking 10 times longer to render.

The Challenge

Improve Rendering Performance and Increase Cluster Utilization

All computer-generated imagery starts with artists who create drawings on paper, developing the concept art that ultimately defines the overall look of the film and the characters. While writers, sculptors, and artists bring a story to life, computers are counted on to perform much of the heavy lifting in animated feature film production. The detail given to shadows and reflections, facial expressions, clothing and background scenery all have one thing in common: they are brought to the screen using complex computer programs.

Sculptors work from artists' drawings to build clay models of the characters in various poses, which are then digitized into 3D wire frame models. Their every movement—from how they blink to how they smile—is articulated and stored on a ModelFarm. Animators work from the data in the ModelFarm to build each frame of film. The rendering is performed in the RenderFarm, a bank of 3200 CPUs that work together to produce the 24 frames needed for each second of film time. It can take several hours to compute one frame of film. One-by-one, the finished frames are stored on the ImageFarm until the entire production is complete and ready for conversion to film.

New Film Techniques Bring New Challenges

Each new film entering production generally requires a much greater magnitude of geometric detail and visual richness, placing massive demands on CPUs and networks.

In this case, the Studio pushed the boundaries of computer animation by incorporating three advanced techniques: 1) irradiance, or the casting of color from one object onto another; 2) ambient occlusion, a technique for making shadows more authentic through the use of multiple light sources; and 3) reflections, made more realistic by taking shape of the object on which they're cast. Considering that there may be as many as 2 billion objects in a shot that need to be rendered, the complexities of embracing these new techniques amidst the backdrop of a 90+-minute feature-length film significantly strained the Studio's file system.

Throughput and Scalability Tax the Legacy Infrastructure

The Studio employed an NFS filer system with NFS caches in front to support the requirements of the six films produced to date. However, once production started on the new film, problems with the RenderFarm quickly became apparent, as frames were taking 10 times longer to render. *(continued)*

“In the early stages of our film, it was taking 10 times longer to render each frame because the NFS caches could not go fast enough to serve up the data needed. Having IBRIX in our network allowed us to get the film done on time—extremely important when you consider the billions in revenue that's at stake from box office receipts and consumer spending on film-related merchandise. Frankly, we weren't going to get there without having a solution that was radically more effective than what we'd used in the past.”

*—Vice President of Technology
Animation Studio*

FAST FACTS

COMPANY

Major California Animation Studio

INDUSTRY/APPLICATION

Media and Entertainment

CHALLENGE

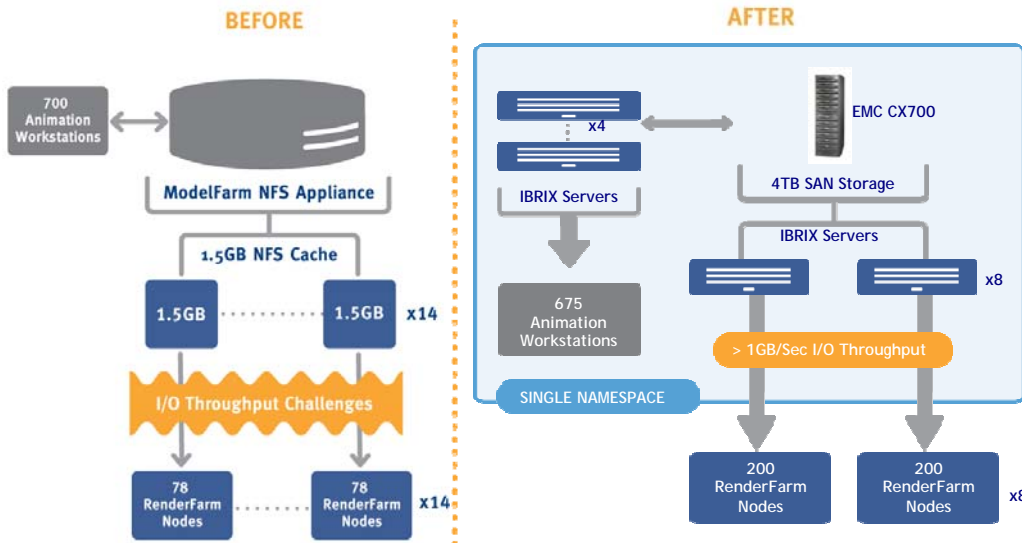
Improve animation rendering time while simultaneously managing peak workloads of up to 240 billion data requests per day by eliminating legacy NFS file system inefficiencies that threatened the Studio's ability to meet its June 2006 release date for their latest feature film.

SOLUTION

IBRIX Fusion™ scalable file serving solution feeding a 3200-processor RenderFarm, integrated with Dell™ PowerEdge™ 1850 and 6850 servers and 4 terabytes of EMC® CLARiiON™ CX700 SAN storage, delivers high throughput performance that allows tremendous scalability, simultaneous access to files, and better balance across multiple nodes, significantly reducing hot spots.

RESULT

The Studio was able to reverse mounting production delays with an incredible 30x increase in I/O throughput that pushed the utilization of their 1000+ node RenderFarm from 18% to more than 80%.



At its peak, the system needed to serve up 240 billion data requests per day. The problem: CPUs in the RenderFarm sat idle while waiting for network I/O because the NFS caches could not go fast enough.

Although bandwidth scalability matters, the scalability of I/O operations per second—and the need to randomly access millions of small files stored in thousands of directories—is paramount. The Studio needed to access 1 billion 4K reads per hour. “It is in small file, metadata-intensive applications that traditional filers and single metadata server-based SAN file systems fall short,” said the Vice President of Technology. “Being able to scale bandwidth independent of capacity is extremely important to us.”

Hotspots proved to be another challenge. To complete each shot, animators worked with a core set of data files, with only about 100 frames active at any given time. The result is that the same data files are served up a million times over, ultimately creating a bottleneck. “Access patterns are unique to each film, and there is no way of predicting usage beforehand. Any solution we chose had to be able to handle our unique

done any faster. Rather it’s a testament to how much more visually rich and computationally complex the films are. Having IBRIX Fusion as our file serving solution is the way to go moving forward.”

Key benefits of the IBRIX solution include:

- Increased Application Performance**
 The Studio realized a 30x increase in I/O throughput that pushed the utilization of the 1000+ node RenderFarm from 18% to more than 80.

- Scalability for Growth**
 With scalability to thousands of nodes and up to 16 petabytes of capacity in a single namespace, IBRIX Fusion eradicates the scalability problems of traditional I/O systems. As a multipurpose and scalable parallel file serving solution, IBRIX Fusion ensures optimal efficiency across a variety of I/O sizes and access patterns and scales equally well and without performance loss for applications characterized by small-block transfers and random I/O.

- Simplified Management**
 IBRIX Fusion distributes metadata information across multiple servers in order to deliver the levels of performance required. All computing elements are managed within a single namespace, providing optimum flexibility for administering the file system and storage assets.

- Reduced Total Cost of Ownership**
 IBRIX Fusion offers the ability to independently scale capacity and performance. “Investing in software capable of running on any hardware is a solution that can be carried forward over time, and it’s what allows us to hit our production deadlines.”

IBRIX solutions continue to be used on projects in production.

The Solution

IBRIX Fusion Scalable File Serving Solution

With the IBRIX solution, the Studio can now scale the infrastructure both ‘up’ and ‘out.’ Vertically scaling performance by adding more commodity cache to the IBRIX I/O servers directly boosts the animation rendering performance. Horizontally scaling performance by adding more IBRIX I/O servers increases I/O throughput linearly from the backend storage to the render nodes. The installation includes IBRIX Fusion installed on 4 Dell™ PowerEdge™ 1850 servers for feeding 675 64-bit Linux animator workstations, and 8 Dell™ PowerEdge™ 6850 servers for processing the data from the ModelFarm to the 1600-node RenderFarm. The solution includes 4 terabytes of SAN storage using EMC® CLARiiON® CX 700 which have been pre-tested and eLab™ certified.

The Result

Performance, Adaptability to Changing Workloads, Long-Term Flexibility

“Our latest film used 300 times more compute power than was used on one of our prior hits eleven years ago, yet each film doesn’t get