

Exclusive for **Siggraph2007**

Solving Rendering Bottlenecks in Computer Animation

A technical overview of centralized storage caching
success in computer animation

> Gear6 accelerates storage and delivers real time application performance with centralized storage caching. Deployed in the data center, Gear6 scalable caching appliances provide scalable and transparent acceleration for existing storage infrastructures to speed up applications, eliminate peak load disruptions, and simplify enterprise configurations.

WHITE PAPER: Solving Rendering Bottlenecks in Computer Animation

INTRODUCTION

Computer generated animation requires enormous amounts of processing and storage horsepower to create dazzling movies. But the production process is always under pressure. Typical movie schedules adhere to fixed release dates coinciding with multi-million dollar marketing launches. In the race to the finish, time saving production techniques provide a chance to improve scene quality, enhance special effects, or spend extra time creating a more engaging storyline. This can lead to bigger box office sales.

While there are numerous steps in creating a digitally animated movie, rendering is a critical element which is both compute and I/O intensive. This paper examines in detail how a centralized storage caching solution accelerates the rendering process to improve time-to-market cycles, streamline storage infrastructure costs, and maximize the productivity of animation teams operating under extreme deadline pressures.

RENDERING OVERVIEW

Movies are made up of scenes, shots, and frames, with rendering taking place at the frame level. Each frame must be generated from a set of models that reside within a single storage repository or model farm. The rendering process is extremely compute intensive because each frame contains multiple models, each of which requires millions of calculations to be properly presented. Similarly, the process is also I/O intensive, as during a frame render, the CPU must have instant access to every model within that frame.

Rendering Performance

Rendering performance is driven by two primary characteristics:

1. Highly parallelizable workload
2. Massive data sharing

Parallelization

Rendering can be easily divided for parallel processing because it represents a collection of discrete, independent units of work such as rendering a single frame. Because individual CPUs process a single frame at a time, adding CPUs shortens the time to render a shot, scene, and ultimately the entire movie. Therefore, animation studios typically acquire rendering farms that include thousands of CPUs.

WHITE PAPER: Solving Rendering Bottlenecks in Computer Animation

Massive Sharing

The parallelization leads to massive data and storage sharing. Specifically, rendering drives sharing of storage resources as each job requires clients to simultaneously access the model farm. With thousands of CPUs each processing individual frames, many of which use the same core models, the impact on the underlying storage system is significant. In particular, disk-based systems have trouble keeping up with the volume and frequency of the requests. The basic architecture of a rendering configuration is shown in Figure 1.

Rendering Clients

- Each render requires simultaneous client access to model farm

Parallelization and massive sharing lead to I/O bottlenecks

Model Farm

- Providing shared I/O to rendering clients




Figure 1: I/O bottlenecks in rendering environments

TRENDS IN COMPUTER ANIMATION PRODUCTION

With a need to bring audiences back for more, movie producers continuously strive to create new and exciting visual effects. This drives two major trends:

Rendering farms are getting bigger

As each production tackles new storylines and detailed special effects, there is a corresponding need to provide the rendering capacity to fulfill the director's vision. For example, water and hair are known to use computationally complex models requiring more CPU cycles for convincing visual representations. As such, animation studios add new CPUs to their render farms with almost every new production.

Data sets are getting larger

Model farms are also increasing with larger data sets. The data sets include larger files based on more complex models, and also more files due to the visual complexity of new releases. For example, scenes with

WHITE PAPER: Solving Rendering Bottlenecks in Computer Animation

large crowds of people or thousands of objects require a larger number of models, and individual models with intricate behavior patterns are greater in size.

MAINTAINING STORAGE PERFORMANCE TO FEED THE RENDER FARM

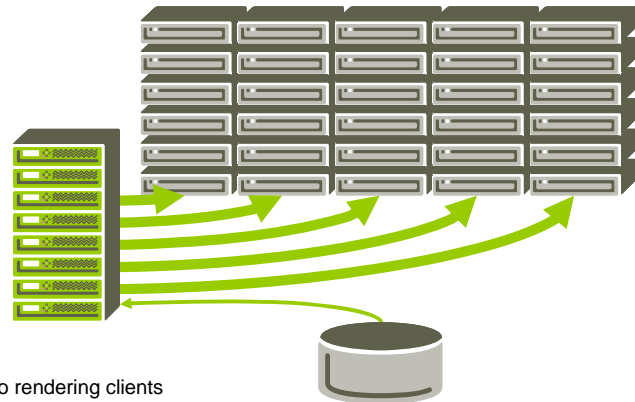
Given the enormous expansion of render and model farms, the goal of the computer animation production system should be to maintain scalability while enabling full utilization of the compute cluster. More specifically, architecture must deliver the following:

1. Scale to accommodate increasing compute clusters (rendering farms)
2. Scale to accommodate increasing data sets (model farms)
3. Meet low latency requirements to maintain 100% CPU utilization
4. Replicate data as appropriate to support massive data sharing

Rendering Clients

- Each render requires simultaneous client access to model farm

Scalable Caching Appliances
Enable full client utilization



Model Farm

- Providing shared I/O to rendering clients

Figure 2: Solving rendering bottlenecks with scalable caching appliances

ARCHITECTURAL OVERVIEW OF SCALABLE CACHING APPLIANCES

Scalable caching appliances from Gear6 address these significant challenges in animation production by delivering transparent and seamless acceleration to existing NFS environments. Deployed as a centralized network resource, scalable caching appliances deliver significant performance improvements

WHITE PAPER: Solving Rendering Bottlenecks in Computer Animation

without requiring a change out of current infrastructure. The following product characteristics outline key architectural advantages of this approach.¹

No new file system

Scalable caching appliances retain and complement existing file systems.

No new persistent storage / no data movement

Scalable caching appliances enhance existing persistent storage instead of replacing it, eliminating rip-and-replace scenarios. Further, achieving a performance boost no longer requires data movement.

Transparency in the data path

Scalable caching appliances deliver a performance boost by accelerating data 10-50 times, but also maintain the ability for customers to access their data directly.

No changes to applications or clients

Scalable caching appliances address storage performance issues by seamlessly applying caching as a shared network resource and using all existing applications "as-is". No changes are required at the client/application or storage management layer. This simplifies deployment, provides maximum investment protection, and adds value without disrupting existing software solutions.

Tune at the appliance level

Scalable caching appliances provide tuning at the appliance level. This delivers more customer control and coordination at the network level to avoid more costly and time consuming changes, upgrades, or modifications at the application or storage levels.

PERFORMANCE ASPECTS OF SCALABLE CACHING APPLIANCES

MEMORY-BASED ACCESS

CACHEfx appliances from Gear6 optimize performance by serving data from memory as opposed to disk, distributing requests in parallel, and responding immediately to each portion of the request.

¹ Adapted from the Gear6 Blog, [Throughput](#), in a post titled "[How About My Snapshots?](#)"

WHITE PAPER: Solving Rendering Bottlenecks in Computer Animation

PARALLEL I/O PROCESSING

CACHEfx appliances support highly-parallel I/O processing for each node within a clustered appliance, providing the ability to receive and respond to I/O requests independently and simultaneously. This generates substantial improvements in I/O throughput compared to traditional disk-based systems that do not operate in parallel and rely solely on slow mechanical disks. A typical I/O sequence is shown in Figure 3.

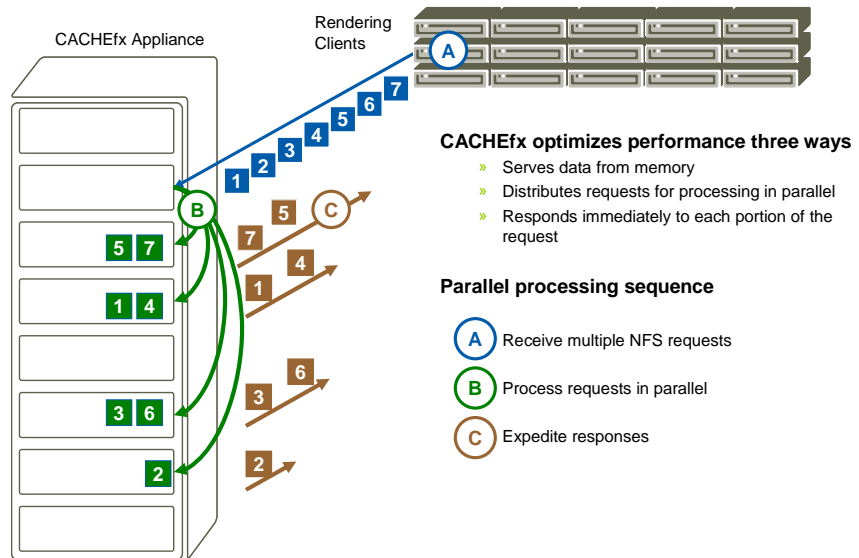


Figure 3: Parallel I/O of a CACHEfx appliance

PERFORMANCE CAPABILITIES FOR RENDERING

Scalable caching appliances deliver critical performance advantages that are ideally suited to rendering environments. These include:

Capability	Rendering Benefit
<ul style="list-style-type: none"> Massive data delivery from a single node Massive data delivery across a single appliance 	<ul style="list-style-type: none"> Scale to accommodate increasing compute clusters (rendering farms)
<ul style="list-style-type: none"> Ability to scale cache capacity to terabyte levels on the fly 	<ul style="list-style-type: none"> Scale to accommodate increasing data sets (model farms)
<ul style="list-style-type: none"> Use high-speed, high-capacity memory efficiently and effectively 	<ul style="list-style-type: none"> Meet low latency requirements to maintain 100% CPU utilization
<ul style="list-style-type: none"> Replicate data as appropriate to deliver massive sharing 	<ul style="list-style-type: none"> Shorten overall rendering time during critical production bottlenecks

WHITE PAPER: Solving Rendering Bottlenecks in Computer Animation

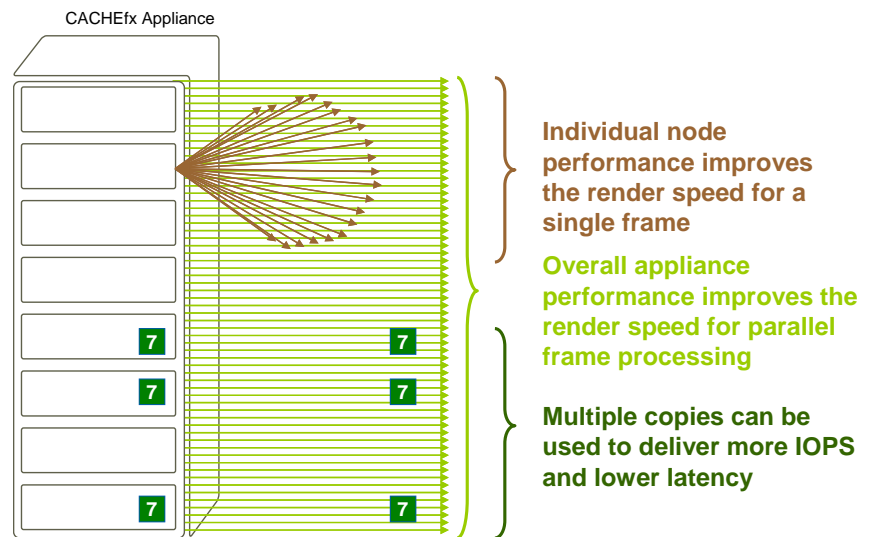


Figure 4: Scalable caching appliance performance capabilities for rendering

PERFORMANCE RESULTS

Figure 5 illustrates CACHEfx capabilities in delivering random access to shared data sets. These benchmarks are based on 1K random reads using between 1 and 7 clients continually issuing random I/O operations. Maximum IOPS exceed 300,000, and latency remains under 250 microseconds even with the addition of new clients.

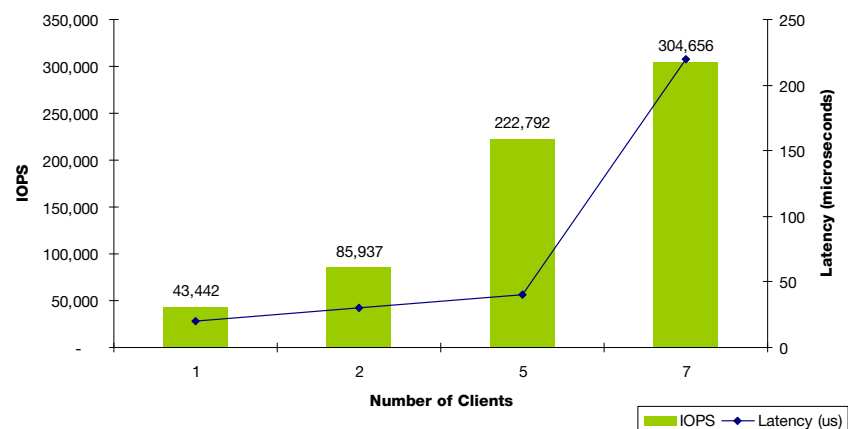


Figure 5: CACHEfx appliance performance results

WHITE PAPER: Solving Rendering Bottlenecks in Computer Animation

CONCLUSIONS

The table stakes in computer animation production are extremely high, given the enormous production costs and ROI expectations of major studios. It is no accident that the animation industry is consistently at the forefront of technology adoption which often leads to multi-billion dollar IT segments. As computer animation continues to be applied to more and more elaborate movie making, the corresponding need to support increasingly sophisticated rendering techniques will require new paradigms to effectively support massive storage infrastructures. Centralized storage caching is being widely adopted within the industry to help push the creative envelope to new extremes.



© 2007 Gear6. All rights reserved. All other trademarks and copyrights are property of their respective owners.
Information provided in this paper may be subject to change.
For more information, visit www.gear6.com