klara

Accelerating Animation Workflows with Optimized ZFS Storage

A case study in high-throughput ZFS design for media workflows

Case Study

Table of Contents

Summary

A leading animation studio partnered with Klara Inc. to deploy a new high-performance storage server for finalized animation projects. Built on OpenZFS, this greenfield deployment was engineered for extreme sequential throughput to handle multi-gigabyte scene files and renders.

By incorporating a carefully designed and tuned ZFS pool layout, and exhaustive hardware validation, the studio achieved 55 GB/s sustained sequential throughput – a 57% increase over the default configuration. This performance boost eliminated render handoff delays and provided headroom for the studio's growing 4K/8K workloads, all while maximizing the value of existing hardware investments.

The Challenge

Early tests showed only 10% of expected bandwidth—revealing that performance required far more than just good hardware.

The animation studio's final render files are enormous, often tens of gigabytes each. Transferring these files from the production storage tier to archive or post-production was becoming a bottleneck.

The studio needed a storage solution that could ingest and deliver multi-GB files at ultra-high speed, ensuring that rendering farms could write out final frames without delay and editors could instantly access completed shots. The goal was to reduce render pipeline delays and accelerate handoff to post-production, keeping artists productive and expensive render nodes busy – not idling while waiting on slow storage transfers. Moreover, any solution had to be future-proof for rising throughput demands (e.g. higher resolution content and more concurrent renders) and delivver strong ROI by leveraging commodity hardware.

Initial evaluations of an out-of-the-box ZFS server showed promise in data management features, but fell far short on raw throughput. In fact, early tests with default settings yielded only a fraction of the hardware's potential – roughly 10% of expected bandwidth in sequential reads/writes. This gap underscored that simply deploying ZFS on high-end hardware was not enough; careful tuning and architectural design would be required to meet the studio's performance targets.

The Solution

Klara's engineering team worked closely with the studio to design and implement a performance-optimized OpenZFS solution tailored to their workflow. The approach spanned from hardware architecture to low-level ZFS parameters:

Optimized ZFS Pool Layout

At the heart of the solution was an OpenZFS storage pool designed for maximum parallel throughput. Klara usually recommends using striped mirror vdevs – essentially creating several redundant pairs of drives, then striping data across those pairs. This layout offers both good data protection and high performance, especially IOPS: each mirror provides fast read access (ZFS can read from both disks in a mirror concurrently) and avoids the penalty of parity calculations.

However, the workload in this case required maximizing throughput without the usual IOPS requirements due to the large file sizes. Utilizing just two RAID-Z2 groups allowed a larger number of NVMe devices to contribute write bandwidth than would have been possible with the striped mirror configuration.

Additionally, with the specific hardware chosen for this deployment, the NVMe drives and 100 GbE network cards were evenly distributed across two NUMA nodes, creating two symmetrical halves of the pool. This architecture ensured balanced load distribution and let the system exploit the full aggregate bandwidth of all devices in parallel.

ZFS Tuning and Configuration

With the hardware and pool in place, Klara fine-tuned OpenZFS settings to match the studio's sequential workload. Key parameters included using a large record size (1 MiB) for the project datasets to optimize for huge sequential reads/writes. By increasing ZFS recordsize from the default 128 KiB to 1 MiB, each I/O operation handled more data, reducing CPU overhead and better

matching the typical sizes of the animation files, while reducing metadata IOPS.

Tuning for raw throughput delivered 55 GB/s performance — a 57% improvement over default ZFS settings.

The team also disabled or adjusted features that, while useful in general, added unnecessary overhead in this scenario. For example, atime (access time) updates were turned off, so the file system wouldn't generate small writes just to note file read timestamps. Deduplication and compression were also kept off for this volume, as the data was already in final compressed formats and the priority was raw throughput over space savings.

Furthermore, the ZFS intent log was tuned for throughput by setting logbias=throughput, telling ZFS to optimize synchronous writes for streaming throughput rather than low latency. These tunings, while subtle individually, together removed many of the bottlenecks and extra work that ZFS normally does – essentially paring ZFS down to a lean, sequential throughput machine tailored for large file streaming.

Validation & Testing

Once the system was configured, rigorous testing was conducted to validate that the delivered hardware met the design goals. The team executed a series of real-world simulation benchmarks, including parallel file transfers and industry I/O tools (e.g. fio with multiple jobs), to measure sustained sequential read and write throughput.

Tests were performed first with default ZFS settings to establish a baseline, then with the optimized configuration. The difference was dramatic: the tuned setup pushed throughput to 55 GB/s on sequential reads (and similarly high on writes), whereas the default config had been much lower.

This 55 GB/s result corresponds to about 440 Gb/s, which in practice means the storage can saturate multiple 200 Gb/s network links for the workflow. In percentage terms, the custom optimizations yielded roughly a 57% improvement in throughput over the baseline out-of-the-box performance. The system scaled well with increased concurrency, thanks to the NUMA-balanced design and removal of bottlenecks, whereas the stock configuration showed signs of CPU saturation and did not improve with additional threads.

Business Impact

The technical gains of the optimized ZFS solution translated directly into business value for the animation studio. With the new storage server in production, the studio experienced several key improvements:

Reduced Render Delays

The high throughput storage eliminated the file transfer bottleneck at the end of the rendering pipeline. Final animation frames and scene files now move off the render cluster 57% faster than before, significantly cutting down idle time. Artists no longer wait hours for large files to copy; instead, renders are finalized and handed off to post-production without delay, keeping projects on schedule

Faster Post-Production Handoffs

The editing and VFX teams received completed animation files sooner, accelerating the downstream processes.

What used to be an overnight transfer of data can now happen in just minutes, enabling near-real-time collaboration between departments. This speed has improved the studio's agility in making last-minute changes and iterating on shots, providing a competitive edge in project delivery timelines.

Maximized Hardware ROI

By unlocking the full performance of their existing hardware, the studio avoided purchasing additional storage systems. The NUMA-aware tuning and ZFS optimizations extracted significant additional throughput from the NVMe drives and network interfaces, delaying any need for costly infrastructure upgrades.

Essentially, the studio achieved tier-1 storage performance using open-source software on standard servers, a testament to high ROI and smart engineering.

Optimized for throughput and balance, the new storage system scales effortlessly with the studio's growing workload.

Future-Proofed Growth

The solution was designed with headroom to spare. At 55 GB/s, the storage system can handle the studio's next few years of workload increases, including the transition to higher resolution content and more concurrent render jobs.

The modular nature of ZFS and the chosen hardware means the studio can scale capacity or performance as needed (by adding more drives or network links, for instance) without a forklift upgrade. This future-ready posture gives the business confidence to take on more ambitious projects knowing the storage backbone will keep up.

What's Next

Buoyed by the success of this deployment, the animation studio is exploring additional ways to leverage OpenZFS across its pipeline.

In the short term, they plan to extend the current solution by adding a second ZFS server at a remote site for replication and disaster recovery – taking advantage of ZFS's built-in snapshot and send/receive features for efficient data sync. The performance headroom also opens the door to consolidating other workloads onto ZFS. The studio is evaluating whether its legacy project archive, currently on an aging NAS, could be migrated to ZFS to benefit from the improved throughput and unified management.

On the technology front, the studio and Klara are keeping an eye on the OpenZFS roadmap. Features like upcoming direct I/O support for cache control and continued NVMe performance enhancements promise even further gains. As those become production-ready, the team intends to incorporate them to stay at the cutting edge of storage speed. The collaboration with Klara continues as a performance and reliability partner for the long term, ensuring the storage infrastructure remains an enabler – not a limiter – of the studio's creative ambitions.

klara

The Klara Inc. team of engineers possesses the deep expertise in the OpenZFS engineering environment required to quickly investigate and rectify even the most stubborn issues. When organizations face technical challenges, they need a partner with experience, industry knowledge, and strong community connections.

klarasystems.com
1.213.634.4466
contact@klarasystems.com

linkedin.com/company/klara-inc
bsky.app/profile/klarainc.bsky.social
facebook.com/klarainc

Looking to learn more?

Our teams are ready to provide you with the details needed to take the next step.