

Solution Brief

High-Performance Media Rendering and Edge Compute Node

Challenges

- High-resolution video (4K/8K), VFX, and real-time rendering require massive GPU resources, often leading to production bottlenecks
- Rendering jobs across multiple nodes or locations can suffer from inconsistent throughput and resource contention
- Large media assets (video, textures, models) must be moved between compute nodes and storage, causing latency and delays

Solution

- High-performance GPU nodes enable faster video rendering (4K/8K, VFX, CGI) and real-time processing to support parallel rendering across multiple nodes
- Integrated storage delivers high throughput for large media files to eliminates bottlenecks between compute and storage
- Enterprise storage provides data protection (RAID, snapshots) to ensure that intellectual property rights are protected at all stages
- It can be easily scaled to support GPU nodes for rendering workloads and expand storage capacity to meet the growing needs of media libraries

QSAN ALTOS an Acer Group Company

Bottlenecks in Media Production and Edge Processing

Media production, content rendering, and edge-based AI workloads are becoming increasingly compute-intensive, latency-sensitive, and distributed. Enterprises and studios face several key challenges:

- **I/O Bottlenecks:** Modern rendering workloads require massive data throughput. Standard network storage often causes "GPU starvation," where powerful GPUs sit idle while waiting for large media assets to load over the network.
- **Storage Scarcity:** Local NVMe storage on compute nodes is fast but limited in capacity. High-resolution projects quickly outgrow local disks, leading to manual data shuffling and increased risk of project delays.
- **Data Fragmentation:** In collaborative environments, creative assets are often scattered across multiple workstations, making it difficult to maintain a single "Source of Truth" for rendering and post-production.
- **Environmental Constraints at the Edge:** Deploying heavy compute at edge locations (such as film sets or remote branch offices) requires a compact yet enterprise-grade setup that can handle continuous high-load operations.

Altos GB10 F1 + QSAN XN1 Integrated Architecture

This solution transforms the rendering workflow by creating a high-speed pipeline between raw asset storage and GPU acceleration.

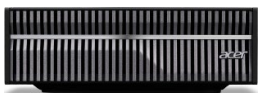
- **High-Throughput Rendering Pipeline:** The Altos GB10 F1 acts as the primary rendering engine, while the QSAN XN1 storage functions as the central asset pool. Connected via 25 GbE with RDMA support, the system allows GPUs to access 8K textures and geometry data directly from the NAS as if they were on local drives.
- **Decoupled Compute and Storage:** By separating compute (GB10) from storage (QSAN), studios can scale each independently. You can add more GB10 nodes to increase rendering speed or expand QSAN capacity to store more projects without disrupting the workflow.

- **Enterprise-Grade Data Resilience:** Unlike local workstation drives, the QSAN XN1 features RAID protection and hardware redundancy. This ensures that even if a drive fails during a rendering job, the data remains intact and the process continues.
- **Edge-Optimized Form Factor:** The compact design of the Altos GB10 F1 and the desktop-friendly QSAN XN1 series make this an ideal "Studio-in-a-Box" for on-site production, DIT (Digital Imaging Technician) carts, and edge AI processing.
- **Streamlined Orchestration:** Utilizing Altos aiWorks 5.0, administrators can monitor GPU health and manage multi-node rendering tasks with ease, ensuring maximum resource utilization.

Best Solution for Media Studios

High-performance GPU rendering combined with edge compute enables faster, low-latency media processing while ensuring secure and scalable content workflows.

Recommended ALTOS Models



Altos BrainSphere™
GB10 F1

- NVIDA® GB10 Grace Blackwell Superchip
- NVIDA® Blackwell GPU 128 GB Unified Memory
- 1 PetaFLOP FP4 AI Computing Performance
- Up to 4 TB NVMe Storage Capacity
- NVIDA® ConnectX®-7 SmartNIC
- One-click AI Developmnet and Deployment Platform

Recommended QSAN Models



QSAN XN1205-4A

- Desktop Tower NAS
- AMD® Ryzen™ 4-core
- 8 GB DDR5 DIMM up to 64 GB
- 3.5” Slot x 4 + 2.5” Slot x 1 + M.2 NVMe x 2
- Maximum Raw Capacity up to 153 TB



QSAN XN1209-4A

- Desktop Tower NAS
- AMD® Ryzen™ 4-core
- 8 GB DDR5 DIMM up to 64 GB
- 3.5” Slot x 8 + 2.5” Slot x 1 + M.2 NVMe x 2
- Maximum Raw Capacity up to 276 TB

QSAN ALTOS an Acer Group Company

© 2026 QSAN Inc. All Rights Reserved. QSAN, the QSAN logo, and the marks listed at www.qsan.com are trademarks of QSAN Inc. Other company and product names may be trademarks of their respective owners.
April 2026