

Solution Brief

Local AI Development and Testing Sandbox

Challenges

- Developers often compete for shared GPU infrastructure, slowing down experimentation and innovation
- Without proper sandboxing, testing new models or code can impact production workloads or other users
- Configuring AI frameworks, dependencies, and toolchains is time-consuming and error-prone
- Difficulty in managing datasets, model versions, and experiment results reduces productivity

Solution

- Developers can quickly spin up isolated environments to enable rapid experimentation without affecting others
- Flexible GPU allocation to maximize utilization while maintaining fairness
- Seamless connection to enterprise storage to eliminate data duplication and transfer overhead
- All development and data remain enterprise infrastructure to ensure data privacy, regulatory compliance, and IP protection

Constraints in AI Research and Development

AI development is inherently iterative, experimental, and resource-intensive, yet many organizations struggle to provide developers with a flexible, secure, and scalable environment. Key challenges include:

- **Prohibitive Cloud Costs:** Continuous R&D, model fine-tuning, and testing on public cloud platforms lead to unpredictable and escalating OpEx (Operational Expenses).
- **Environment Inconsistency:** AI developers often struggle with "dependency hell", where code works on a local laptop but fails in production due to inconsistent software stacks or driver versions.
- **Data Volume and Versioning:** Modern AI requires massive datasets and numerous model versions. Managing these on local workstation drives leads to storage exhaustion and difficulty in tracking data lineage.
- **Collaboration Silos:** Without a centralized repository, sharing large datasets or pre-trained weights between team members is slow and inefficient, often relying on physical drives or slow network transfers.

Altos GB10 F1 + QSAN XN1 Development Sandbox

This integrated solution provides a unified platform that accelerates the "Code-Train-Test" cycle.

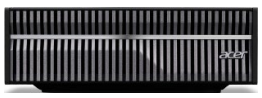
- **Centralized Data and Image Repository:** The QSAN XN1 storage acts as the primary vault for all AI assets. Developers can store large-scale datasets, Docker images, and multi-gigabyte model weights (checkpoints) in a single, RAID-protected location accessible to the entire team.
- **High-Performance Compute Node:** The Altos GB10 F1, equipped with NVIDIA GPUs, provides the raw power needed for heavy training and inference testing. Its architecture is optimized for high-density compute, ensuring that developers spend less time waiting for jobs to finish.

- **Rapid Workspace Provisioning:** Utilizing Altos aiWorks 5.0, teams can quickly deploy standardized development environments via containers. These containers can directly mount high-speed volumes from the QSAN XN1 storage via RDMA, ensuring that data-heavy training sessions are not bottlenecked by network speed.
- **Snapshot-Based Experimentation:** Leveraging the QSAN snapshot technology, developers can "freeze" a dataset or a model environment before an experiment. If the data becomes corrupted or the test fails, the environment can be restored to its original state in seconds.
- **Secure Sandbox Isolation:** The entire stack resides within the enterprise LAN (Local Area Network). This allows developers to test models on sensitive, real-world corporate data without the risk of external leakage.

On-Prem AI Development Environment

An on-prem AI sandbox empowers developers with secure, on-demand GPU environments for rapid experimentation and scalable model development.

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