

Solution Brief

# Agentic Intelligent Document Processing with QSAN XN1 and Altos GB10

## Challenges

- AI-generated outputs often lack clear linkage to source documents or evidence
- Enterprises cannot rely on outputs that are not verifiable or compliant with regulations
- Document ingestion, AI inference, and result storage are disconnected across systems
- Multi-model workloads (LLM + vision + parsing) require stable, high-throughput infrastructure

## Solution

- Multi-step AI processing enables intelligent, context-aware automation
- Every output is linked to ensures transparency and trust
- All data stored on QSAN XN1 provides centralized data control, versioning and audit trails, secure enterprise storage
- Powered by Altos GB10 supports multi-model pipelines (LLM + vision) and optimized for batch and continuous workloads to delivers consistent performance at scale

## The Gap Between GenAI and Business Reality

Enterprise data is predominantly unstructured documents (PDFs, images, contracts, forms), but traditional AI and GenAI solutions face critical limitations:

- **Unstructured Data Overload:** The vast majority of enterprise data is trapped in PDFs, images, and manual records that traditional OCR cannot contextually understand.
- **Lack of Traceability:** Generic GenAI tools often provide answers without evidence, making them unsuitable for regulated industries (finance, legal, government) where results must be auditable.
- **Security & Compliance:** Sending sensitive internal documents to public cloud AI models risks data leakage and violates strict data sovereignty requirements.

## Altos GB10 F1 + QSAN XN1 Integration

Altos delivers a dedicated on-premise AI appliance designed specifically for document-heavy workloads, integrating LLM (Large Language Model), Vision, and Document Parsing into a single pipeline.

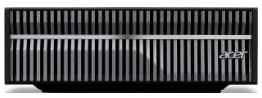
- **End-to-End Automated Pipeline:**
  1. **Input:** Documents are ingested directly from the QSAN X1 storage.
  2. **Processing:** The Altos GB10 F1 runs Agentic AI workflows to parse, analyze, and extract data.
  3. **Verification:** The system generates not just the "answer," but the "evidence" (citations from the source document) to ensure auditability.
  4. **Output:** Results and structured data are saved back to the QSAN XN1 storage for downstream business use.

- **Optimized for High Throughput:** Unlike generic servers, this architecture is optimized for batch processing and stable continuous workloads, capable of handling massive document volumes without performance degradation.
- **Data Sovereignty:** By keeping the entire compute and storage stack on-premise, no data ever leaves the enterprise boundary, ensuring 100% compliance with local regulations.
- **Ready-to-Deploy Edge AI:** Features NemoClaw / OpenClaw integration on the GB10, providing a mature path for edge-AI development and immediate business deployment.

## NAS-Aware Agentic IDP Architecture

Enterprises are overwhelmed by unstructured data that remains siloed in documents. The Agentic IDP solution leverages the power of the Altos GB10 F1 and QSAN XN1 to transform static documents into traceable, actionable business intelligence through a high-throughput, NAS-aware AI pipeline.

## 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