OLTAIX™

THE AGENTIC AI-POWERED PLATFORM FOR ASSET OWNERS™

An Institutional AI White Paper

Executive Summary

Institutional Asset Owners (AOs)—pensions, sovereign wealth funds, insurers, and endowments—steward some of the world's largest pools of capital. Yet their ability to govern effectively is constrained by structural fragmentation: custodians, asset managers, and external data providers each deliver siloed reports that rarely align.

The result is delayed reconciliations, inconsistent valuations, operational inefficiency, and governance risk. Boards and trustees often find themselves looking backward at static reports rather than forward with actionable foresight.

OLTAIX™ is the first AI-native platform built exclusively for Asset Owners. It delivers sovereign, conflict-free intelligence by unifying inputs across servicers, managers, and providers—while respecting each partner's autonomy. The outcome is a single source of evidence-backed truth that equips boards, CIOs, and trustees to govern with clarity, confidence, and foresight.

The Business Problem

- ❖ Fragmented Data Across Partners - Custodians, managers, administrators, and providers deliver siloed feeds. Without integration, there is no single, real-time, authoritative source of truth.

- ❖ Backward-Looking Reporting - Legacy systems emphasize static, historical reporting. Trustees and boards are left reacting to the past rather than anticipating risks and opportunities.

- ❖ Private Markets Complexity - Valuations are opaque and delayed. Capital calls and distributions are difficult to align with liquidity forecasts, eroding confidence in allocation decisions.

- ❖ Operational Inefficiency - Exception management, reconciliations, and compliance checks remain highly manual and resource-intensive. Cycle times slow decisions and raise costs.

- ❖ Regulatory & Fiduciary Pressure - Regulators demand greater transparency into liquidity, risk, fees, and social-impact metrics. Trustees must demonstrate evidence for every decision—yet legacy platforms struggle to deliver.

- ❖ Fee Transparency & Oversight - Scrutiny of investment and servicing fees is intensifying. Current systems make it difficult to attribute costs across managers,

mandates, and vehicles.

- ❖ Governance Burden - Boards receive large volumes of disconnected reports, forcing trustees to manually piece together insights. Lack of explainability undermines fiduciary confidence.

- ❖ Technology Debt - Most platforms are outdated architectures never designed for AI or modern data demands. Customization is slow and costly, leaving AOs dependent on vendor roadmaps.

- ❖ Risk & Liquidity Management Gaps - Stress testing, liquidity ladders, and risk attribution often lag reality. During market shocks, decision-makers lack the real-time foresight they need.

- ❖ Talent & Knowledge Constraints - Experienced staff are retiring, and younger talent is scarce. Manual, repetitive work discourages retention and risks knowledge loss.

How OLTAIX Addresses These Challenges

- ✓ Unified Intelligence Plane - Brings all partner feeds into a single AO-controlled layer without disrupting partner autonomy.

- ✓ Predictive Foresight - LLMs and RAG deliver forward-looking insights with evidence citations, enabling boards to see what's next.

- ✓ Private Market Transparency - Dedicated agents reconcile NAVs, capital calls, and distributions for liquidity alignment.

- ✓ Automated Workflows - Agentic AI streamlines reconciliations, compliance, and exceptions, cutting cycle times significantly and freeing staff for higher-value work.

- ✓ Auditable Governance - Every output is tied to underlying data, producing a complete fiduciary audit trail.

- ✓ Real-Time Fee Attribution - Agents track and reconcile fees across mandates, providing clarity on true costs.

- ✓ Board-Ready Narratives - Reports are auto-generated in plain language with citations, reducing reporting burdens.

✓ AI-Native Architecture - Purpose-built with agentic AI, MCP, and RAG for scalability and customization.

✓ Continuous Risk Oversight - Foresight agents run stress tests and liquidity ladders in real time.

✓ Cognitive Co-Pilot - Automates repetitive work, preserves institutional knowledge, and augments scarce talent.

Defining the AI Technologies Involved

• Large Language Models (LLMs): The "reasoning brain" of the platform—able to read documents, explain complex data, and produce board-ready narratives.

• Retrieval-Augmented Generation (RAG): Ensures AI insights are grounded in actual AO data, always evidence-backed and cited.

• Model Context Protocol (MCP): A secure standard that lets AI safely call tools, systems, and data sources without compromising governance.

• Agentic AI: Multiple specialized agents—reconciliation, compliance, attribution—that collaborate, cross-check, and deliver auditable results.

The OLTAIX™ Solution — AI-POWERED FORESIGHT FOR ASSET OWNERS™

OLTAIX™ provides Asset Owners (AO) with sovereign control, foresight, and clarity. Its architecture allows each partner type—Asset Servicers (AS), Asset Managers (AM), and External Data Providers (DP)—to maintain its own agentic cluster, while the AO integrates all insights through the OLTAIX Intelligence Plane.

The AS1…N / AM1…N / DP1…N topology makes the model inherently scalable: every new partner is simply onboarded as another autonomous cluster, ensuring customization without complexity.

Topology Illustration

The following diagram illustrates an illustration only of a potential OLTAIX™ partner topology. A more detailed topology, with **unlimited** number of partners (1 to N) is included in the **Appendix.**
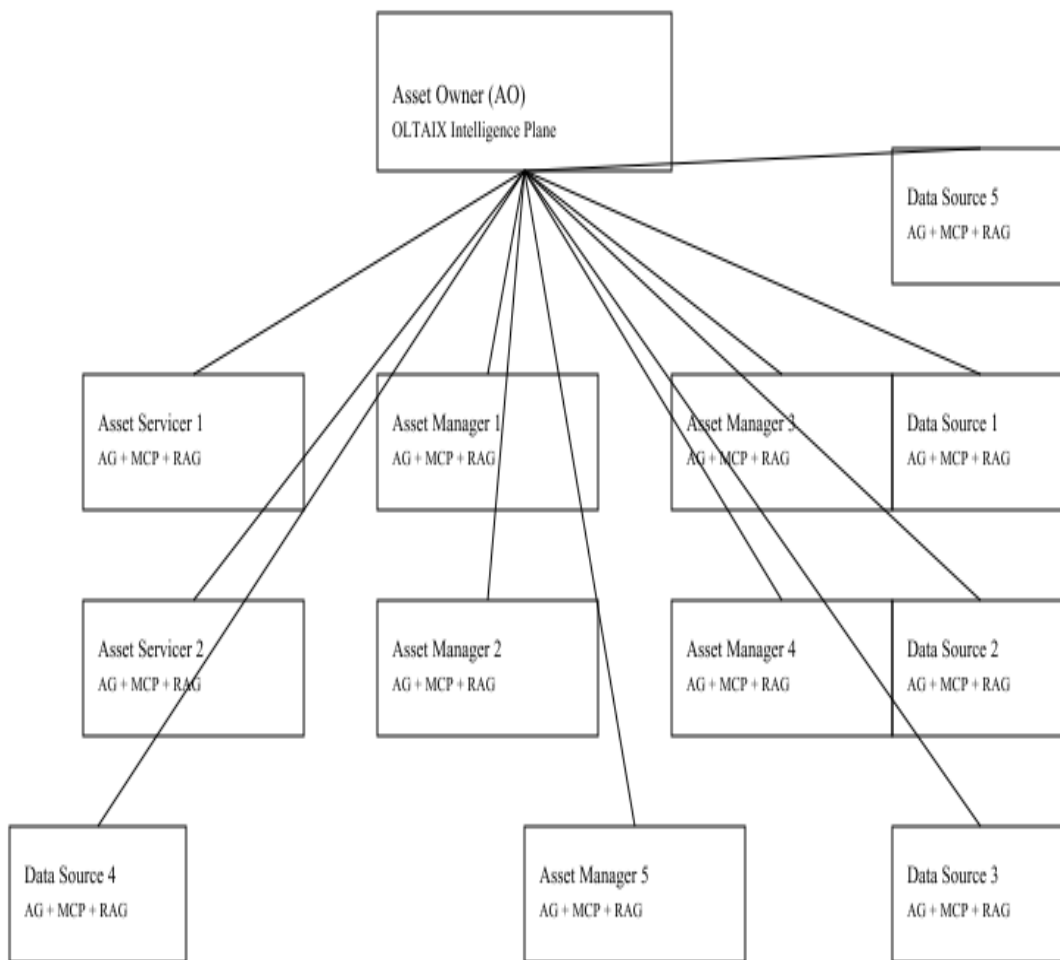
Why OLTAIX™ is Unique

• Exclusively for Asset Owners: Conflict-free, independent of asset managers.
• AI-Native by Design: Built ground-up, not retrofitted, for institutional governance and scale.
• Customizable: Each AO gets its own sovereign intelligence plane, tuned to its ecosystem.
• Evidence-Backed: Every output is cited, auditable, and ready for fiduciary review.
• Scalable: Whether 2 servicers today or 10 tomorrow, OLTAIX scales linearly with partner growth.

Workflow Example: T+1 NAV Variance

1. AO detects a NAV variance >50bps.
2. Servicers review reconciliations; Managers confirm exposures; Providers check index events.
3. Each AI cluster returns evidence-backed results.
4. AO consolidates findings into its Unified Intelligence Plane, identifies the root cause, and proposes remediation.
5. Compliance approves dual-control execution; corrections flow back to partners.

Governance & Security

• Sovereign Control: AO owns orchestration, policy, and intelligence.
• Zero-Trust Boundaries: Partners remain isolated with least-privilege access.
• Data Minimization: Only necessary data crosses trust boundaries; sensitive data masked.
• Model Governance: Versioned prompts, evals, drift monitoring, and audit logs.
• Evidence-Based: No black boxes—every output is linked to its sources.

Conclusion

OLTAIX™ redefines institutional intelligence. By addressing the ten core challenges that Asset Owners face, it delivers sovereign, conflict-free, evidence-backed foresight. Its customizable, scalable topology ensures every AO can govern with confidence, independence, and clarity.

Disclaimer and Forward-Looking Statements

This document contains forward-looking statements regarding OLTAIX™, Institutional AI, and the potential use of agentic AI technologies within the institutional asset owner ecosystem. These statements are based on current expectations, assumptions, and projections and are subject to known and unknown risks, uncertainties, and other factors that may cause actual results, performance, or developments to differ materially from those expressed or implied.

OLTAIX™ is not a pre-packaged product. It is presented as a design vision and customizable framework to be tailored for each institutional asset owner. Every deployment is expected to be designed around the specific custodians, managers, data providers, and governance needs of the client. Features, functions, and capabilities are client-specific, subject to ongoing design, and may evolve over time.

Nothing in this document should be construed as a guarantee, commitment, or representation of future functionality, performance, delivery, or outcomes. Institutional AI undertakes no obligation to update or revise forward-looking statements in response to new information, future events, or otherwise.

This document is provided strictly for informational purposes only. It does not constitute investment, legal, regulatory, or other professional advice. Institutional AI makes no warranties, express or implied, regarding the accuracy, completeness, or suitability of the information contained herein. Any reliance on the content of this document is at the sole discretion and risk of the reader.
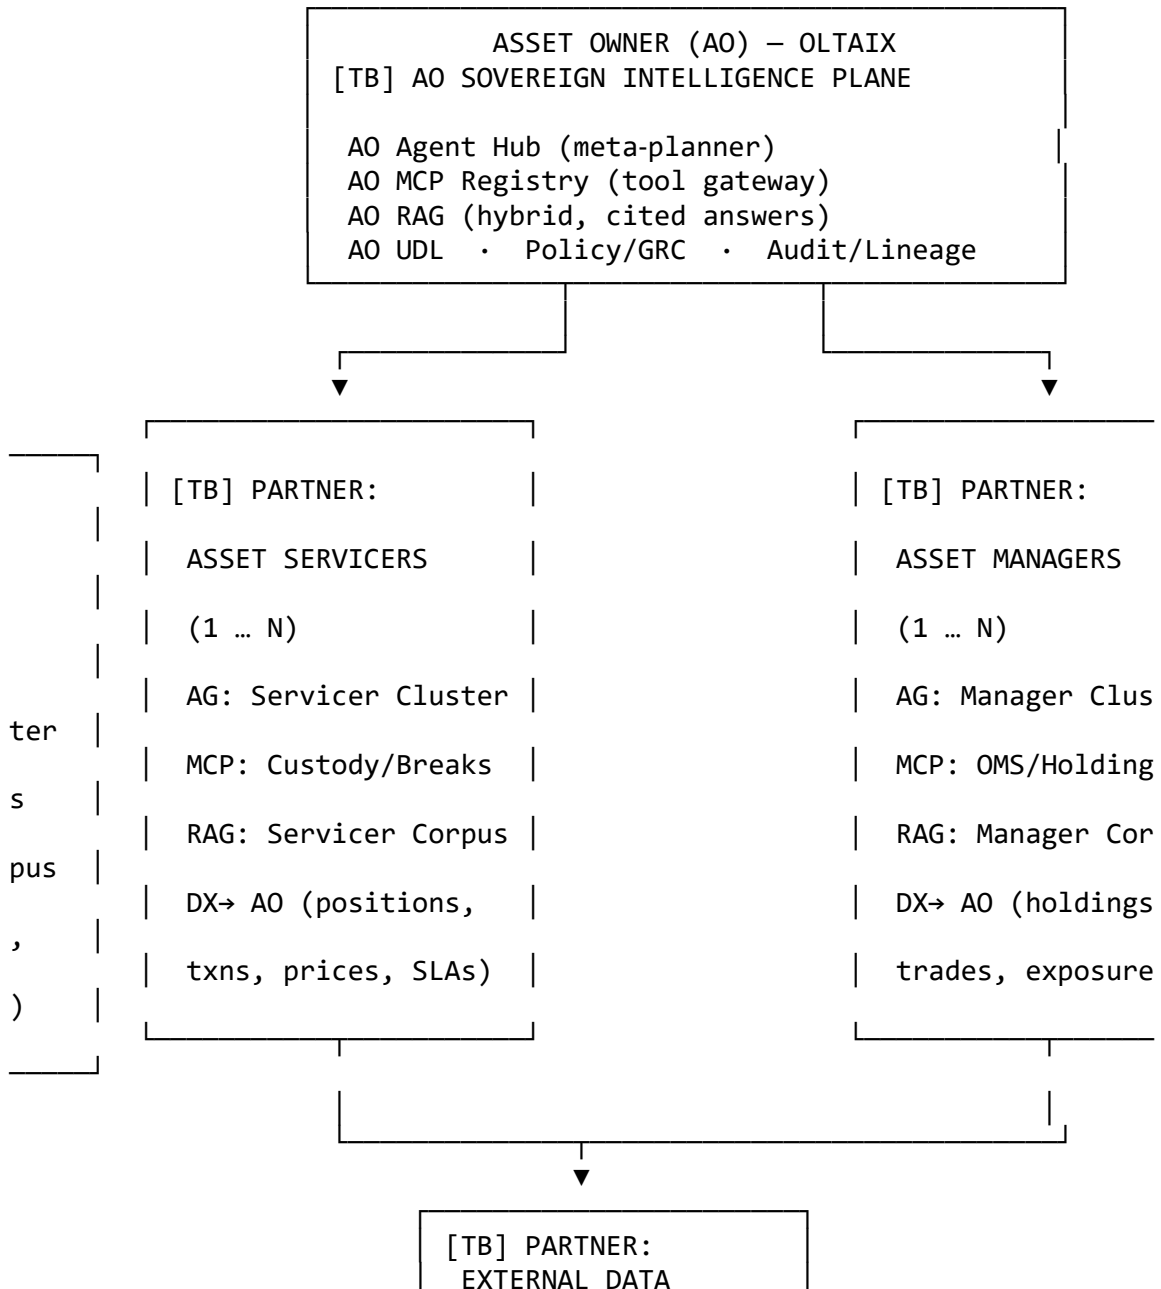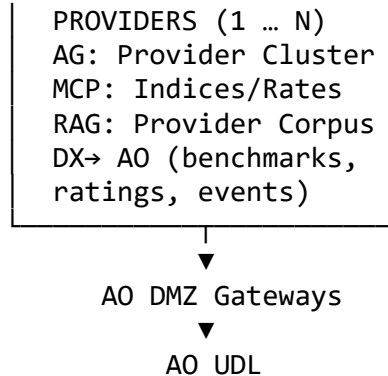
APPENDIX

**OLTAIX™ TOPOLOGY — ONE ASSET OWNER WITH THREE PARTNER TYPES,
WITH UNLIMITED NUMBER OF PARTNERS IN EACH CATEGORY (1 to N)**

Goal: One Asset Owner (AO) operating OLTAIX where each **partner type** runs its own **Agentic AI cluster (AG)** with a dedicated **MCP registry** and **RAG index**. AO keeps sovereign control, evidence, and audit.

Legend: [TB]=Trust Boundary · AG=Agentic AI (planner/executor/critic) · MCP=Model Context Protocol · RAG=Retrieval-Aug. Gen. · DX=Data Exchange · UDL=Unified Data Lakehouse · GRC=Governance/Risk/Compliance

1) High-Level Topology (compact ASCII for Word portrait)

```
                ASSET OWNER (AO) — OLTAIX
         [TB] AO SOVEREIGN INTELLIGENCE PLANE

          AO Agent Hub (meta-planner)
          AO MCP Registry (tool gateway)
          AO RAG (hybrid, cited answers)
          AO UDL   ·   Policy/GRC   ·   Audit/Lineage


                           ▼                              ▼

            | [TB] PARTNER:     |        | [TB] PARTNER:
     |      |                   |        |
     |      |  ASSET SERVICERS  |        |  ASSET MANAGERS
     |      |                   |        |
     |      |  (1 … N)          |        |  (1 … N)
     |      |                   |        |
ter  |      |  AG: Servicer Cluster |    |  AG: Manager Clus
s    |      |                   |        |
pus  |      |  MCP: Custody/Breaks  |    |  MCP: OMS/Holding
     |      |                   |        |
,    |      |  RAG: Servicer Corpus |    |  RAG: Manager Cor
)    |      |                   |        |
            |  DX→ AO (positions,   |    |  DX→ AO (holdings
            |                   |        |
            |  txns, prices, SLAs)  |    |  trades, exposure


                           ▼

                  [TB] PARTNER:
                   EXTERNAL DATA
```

```
┌─────────────────────────┐
│ PROVIDERS (1 … N)       │
│ AG: Provider Cluster    │
│ MCP: Indices/Rates      │
│ RAG: Provider Corpus    │
│ DX→ AO (benchmarks,     │
│ ratings, events)        │
└─────────────────────────┘
            │
            ▼
     AO DMZ Gateways
            ▼
        AO UDL
```

2) Design Principles

1) **Sovereign AO Control:** AO owns meta-planner, unified MCP gateway, authoritative RAG, and UDL. Partners cannot see each other.
2) **Per-Type Autonomy:** Each partner **type** (Servicers, Managers, Providers) operates its own AG + MCP + RAG; scope limited to its artifacts/APIs.
3) **Contracted DX:** Typed feeds/events cross boundaries via schema contracts, CDC/eventing, and policy enforcement (minimize data, mask PII).
4) **Trust Boundaries:** All calls traverse AO DMZ with allow-lists, rate limits, content filters, and entitlements-aware retrieval.
5) **Cite-or-Fail & Audit:** Every output is evidence-backed with citations and full prompt/action lineage.

3) What Lives Where

Asset Owner (AO)

- **AO Agent Hub:** Orchestrates multi-party workflows (recon, valuation, compliance).
- **AO MCP Registry:** AO-approved tools + partner gateway adapters.
- **AO RAG:** Consolidated, deduped, point-in-time corpus (positions, cash, breaks, capital calls, policies, contracts, minutes).
- **AO UDL:** Versioned facts + semantic layer; **GRC** for IPS/limits, liquidity ladders; **Audit** for prompts/tools/actions.

Asset Servicers (1 … N)

- **AG:** Exceptions detection, reconciliation, SLA status, remediation playbooks.
- **MCP:** Custody APIs, exceptions DB, corp actions, pricing files.
- **RAG:** Procedures, file specs, KBs, ticket history.
- **DX→ AO:** Positions, transactions, prices, breaks, timeliness events.

Asset Managers (1 … N)

- **AG:** Mandate compliance, exposure deltas, performance commentary.
- **MCP:** OMS/EMS, holdings/exposure APIs, fee schedules.
- **RAG:** PM letters, guidelines, model books, factor notes.
- **DX→ AO:** Holdings, trades, exposure, commentary, accruals.

External Data Providers (1 … N)

- **AG:** Index changes, ratings actions, macro series, social-impact metrics.
- **MCP:** REST/FTP/S3 access, entitlements checks, throttling.
- **RAG:** Methodologies, licenses, coverage notes, point-in-time guides.
- **DX→ AO:** Benchmarks, classifications, curves, events (rebalances/downgrades).

---

4) Orchestration Pattern (Example: T+1 NAV Variance > 50 bps)

1) **Plan:** AO meta-planner creates sub-goals: Servicers (breaks/SLAs), Managers (exposure/factors), Providers (index events).
2) **Delegate:** AO calls each partner AG via MCP gateway; each performs local RAG retrieval.
3) **Gather:** Partners return cited artifacts (exceptions file, trade/exposure notes, index bulletin).
4) **Reason & Propose:** AO aligns timelines in UDL, explains root cause, proposes actions with confidence scores.
5) **Approve & Act:** Dual-control actions; tasks opened; DX requests issued; everything logged.

---

5) Security, Governance & KPIs

- **Zero-Trust:** Token-scoped, least-privilege, time-boxed access across [TB]s.
- **Data Minimization:** Only necessary fields/rows cross boundaries; masking/redaction at ingest/serve.
- **Model Governance:** Versioned prompts, eval suites, drift monitors, cost/latency SLOs.
- **KPIs:** Evidence coverage ≥95%; query-to-cited-answer P95 <120s; −60% exception cycle time by Day 90.

Summary: Each partner type runs its own autonomous Agentic stack (AG+MCP+RAG). The AO retains sovereign intelligence, merges evidence, and executes auditable, decision-grade workflows in OLTAIX.