Universal Data Bridge

Investor Whitepaper — Unlocking Trapped Value in Legacy Systems

Version 1.0 — January 2026

1. Executive Summary

Across logistics, manufacturing, healthcare, and finance, core operations still run on legacy systems that are reliable but disconnected from modern cloud tooling. The result is expensive manual “swivel-chair” work: humans copying data from documents and screens into other systems. Universal Data Bridge eliminates this “air gap” using a non-invasive approach: Vision AI plus deterministic UI automation, with built-in human review and audit trails.

2. Problem: The Operational Air Gap

Legacy applications frequently lack APIs, modern authentication, or integration paths. Replatforming is slow and high-risk. Meanwhile the business requires near-real-time reporting, customer visibility, and automation. Organisations compensate by adding people — a scalable cost centre that also introduces errors and compliance incidents.

- High ongoing labour cost for data entry and reconciliation
- Error rates compound into chargebacks, shipping mistakes, and financial restatements
- IT reluctance to modify legacy code leads to integration paralysis

3. Solution: Non-Invasive UI-Layer Integration

Universal Data Bridge operates at the UI layer. Agents observe screens and documents, extract structured data, and enter information into target systems at safe human-like speeds. This bypasses the need for APIs, database access, or vendor-led upgrades.

The loop: Capture → Extract → Validate → Review (optional) → Act → Audit.

4. Product Architecture

The product is split into two components: (1) a Next.js control plane for demos, review workflows, and operator oversight; (2) a Node/TypeScript engine using Playwright for deterministic UI control and a multimodal model for vision extraction. Workflows are modular and can be produced quickly using the recorder and generator.

Layer	Purpose	Why it matters
Control Plane (Next.js)	Review UI, demo UI, APIs for approval	Enables human-in-the-loop governance and auditabilit
Automation Engine (Node)	Browser automation, workflow runner, recorder	Deterministic, repeatable execution
Vision Extraction	Multimodal extraction into JSON w/ confidence	Robust against layout drift vs classic OCR
Audit & Evidence	Structured logs and screen captures	Procurement, compliance, and troubleshooting

5. Differentiation

Traditional RPA tools are coordinate-driven and brittle. Classic OCR is fragile when scans degrade or layouts shift. Universal Data Bridge differentiates by combining: (a) vision extraction that returns structured objects with confidence and evidence, and (b) deterministic automation that is versioned, replayable, and governed.

- Concept-aware interactions (control identification beyond fixed coordinates)
- Human Review Mode (operator approval as a default safety valve)
- Workflow Recorder + Generator (fast onboarding of new processes)
- Audit-grade evidence trail (events + screenshots)

6. Go-to-Market

The primary wedge is operations teams trapped in manual rekeying: AP invoice entry, inventory updates, claims intake, shipping label creation, or customer onboarding. A 30-day pilot with clearly defined throughput and accuracy metrics converts into recurring bot-seat licensing.

Target buyers: Head of Ops, Finance Ops, Shared Services, IT integration leads, and compliance stakeholders.

7. Unit Economics and Pricing Rationale

Pricing is anchored to labour displacement and error reduction. A single bot seat typically replaces 0.5–2 FTE-equivalents of repetitive data handling, depending on process complexity and volume. Margin profile improves with workflow reuse and recorder-driven onboarding.

8. Roadmap (12–24 months)

- Confidence-based retries and multi-pass capture for noisy UIs/documents
- Client configuration packs (selectors, validation rules, retention policies)
- Enterprise connectors (SIEM export, SSO/RBAC, secrets management)
- On-prem/VPC deployment with controlled model routing
- Governed workflow marketplace for reusable patterns

9. Risk and Mitigations

- **Extraction uncertainty:** mitigated via confidence scoring + mandatory review for low confidence.
- **UI drift:** mitigated via recorder refresh, selector strategies, and detection tests.
- **Security:** least privilege, local evidence storage options, and enterprise deployment modes.
- **Compliance:** auditable decision trail; separation of extraction and execution.

Appendix: Demo Walkthrough

The included repository contains a live demo legacy UI at /demo/legacy and a web review UI at /review. Run the engine with --review-web to produce a pending extraction, approve it in the browser, then observe automated entry.