

REVENUE OPERATIONS

HERALD.

Win more. Deliver faster.

The capture-to-delivery operations platform for organizations responding to high-stakes solicitations. Federal contractors. State and local procurement. Commercial RFP responses. Healthcare and education contracting. One platform from the solicitation drop through orchestrated post-award delivery, regardless of which procurement system the buyer uses.

BUILT FOR Proposal directors and capture managers across federal, state, and commercial markets · Federal sales leaders (FedCiv, DoD, IC) · State and local procurement response teams · Commercial proposal teams responding to enterprise RFPs · Contracts and compliance officers

WHY NOW

Adoption of generative AI among proposal teams doubled in twelve months, from 34 percent to 68 percent. The firms compounding that advantage are the ones running capture, proposal, contract assembly, and post-award workflow on a single platform, not bouncing between three to five disconnected tools and not limited to one procurement market.

CORE FEATURES

Intelligent Solicitation Decomposition

Modern solicitations are dense and unforgiving. HERALD automatically parses any structured solicitation, federal Section L and M, state procurement codes, commercial RFP structures, healthcare and education ITBs, mapping every requirement to its proposal location and proving traceability line by line. Every reviewer question has a documented answer before the question is asked.

Procurement Portal Orchestration

Procurement portals consume hours of senior capture time on data entry that adds no value to the proposal itself. HERALD populates federal portals (SAM.gov, eBuy), state and local portals, and commercial vendor systems directly from the proposal record, eliminating duplicate entry and the formatting errors that come with it.

Award-to-Delivery Continuity

Most teams treat contract assembly as a separate project from proposal development, losing weeks of momentum after award. HERALD assembles MSAs, SOWs, Task Orders, MAS schedules, GWAC variants, and commercial contract templates from the awarded proposal, so delivery starts the day after award rather than three weeks later.

Color-Team Workflow Orchestration

Pink, Red, and Gold team reviews, sign-off chains, and post-award onboarding all live in the same record. The federal proposal cadence and commercial gate review structures both fit. The proposal becomes the contract becomes the delivery plan, with no information loss between phases.

Multi-Market Compliance Posture

FAR, DFARS, CMMC, FedRAMP, ITAR, and EAR for federal work. State procurement codes and cooperative purchasing terms for state and local. SOC 2, ISO 27001, and HIPAA for commercial and healthcare contracting. The right clauses surface inline as proposals are drafted, by market.

Adaptive Drafting Intelligence

Trained on the firm's own proposal library and prior wins, the platform learns the cadence, evidence patterns, and voice that produce winning narratives. It does not write generic boilerplate; it drafts the way the team's best proposals are written, regardless of which procurement market the response targets.

KEY USE CASES

Multi-Agency Federal Pursuit

Coordinating capture across FedCiv, DoD, and IC opportunities with different compliance regimes, evaluation criteria, and submission portals from a single workspace.

State and Local Procurement Volume

High-frequency state and local RFP responses where the team needs federal-quality compliance discipline without federal-only tooling. Cooperative purchasing terms, state procurement codes, and per-state portals handled inline.

Commercial RFP Response Programs

Large enterprise commercial RFPs in SaaS, professional services, healthcare, and education where the volume and stakes justify dedicated tooling rather than spreadsheet-driven response work.

GWAC and MAS Task Order Volume

High-frequency task order issuance under existing vehicles where speed and compliance are the binding constraints on win rate.

Capture-to-Contract Continuity

Eliminating the post-award scramble by converting proposal commitments directly into contract clauses and delivery obligations, federal or commercial.

ENTERPRISE QUALITY

VALIDATE.

Patch with confidence.

The enterprise validation platform that builds and maintains a living map of the client's entire system, then autonomously verifies that the system continues to work as the business expects it to work, release after release.

BUILT FOR IT and application leadership responsible for production system stability · QA leads supporting ERP, HCM, finance, and commerce environments · Patch and release managers under quarterly update cadence · Compliance teams documenting patch impact and access changes

WHY NOW

Manual regression testing of a single quarterly patch consumes hundreds of person-hours and produces a spreadsheet leadership cannot interpret. VALIDATE compresses the work to hours and produces a leadership-ready readout, so go-live decisions are made on evidence rather than hope.

CORE FEATURES

Adaptive Self-Healing UI

Traditional testing breaks when a screen changes. VALIDATE intuitively recognizes interfaces, adapting to layout shifts and UI updates in real-time. By self-correcting its own pathways, it ensures uninterrupted validation and completely eliminates test maintenance debt.

Dynamic System Library

Rather than relying on rigid, manually authored test scripts, VALIDATE builds and maintains a living map of the client's entire system. It autonomously inventories screens, workflows, and interactive elements, meaning it understands the environment out-of-the-box and grows smarter with every execution.

True User-Journey Validation

The platform does not just ping servers in the background. It autonomously navigates and visually verifies the actual human workflow. It confirms that the right data appears on the right screen exactly as a human user would experience it.

Executive-Ready Translation

Stripping away dense system jargon, the platform synthesizes complex system behaviors into clean, plain-English impact reports. It bridges the gap between system execution and human decision-making, instantly generating the actionable insights leadership needs to authorize a go-live.

Predictive Risk Mapping

By analyzing patterns across multiple runs, VALIDATE acts as a proactive safeguard. It identifies recurring bottlenecks, system degradation, and high-risk workflows, allowing teams to address vulnerabilities before they impact live operations.

Frictionless Enterprise Alignment

VALIDATE is entirely agnostic to the underlying software. Whether a client is running finance, HR, or commerce systems, the platform seamlessly maps to their unique business pathways to provide universal, end-to-end assurance.

KEY USE CASES

Security Profile Changes

Managing and validating access impacts across multiple business units, departments, and areas of responsibility.

Report Modifications

Validating data accuracy post-modification and preventing security lapses across HCM, finance, and operational reporting.

Implementation and Migration Testing

Validating impacts from new module additions, setup changes, release upgrades during implementation, and ensuring stable configurations when migrating to new test or development environments.

Routine Maintenance

Seamless regression testing for quarterly patches, UI updates, and ongoing UX research.

WORKFORCE INTELLIGENCE

SCOUT.

First-round hiring, end-to-end. Compliance-native.

The AI recruiting platform that handles the entire first-round hiring funnel. JD pressure-testing, resume scoring with contact extraction, mandatory screening, AI-assisted interviewing with per-turn AI-use detection, continuous adverse-impact analytics, and audit-ready compliance documentation. One platform, one defensible workflow, from job posting to qualified shortlist.

BUILT FOR Heads of Talent Acquisition and Recruiting Operations · Chief People Officers · Hiring teams running high-volume requisitions · Hiring compliance and employment law counsel

WHY NOW

Two shifts in hiring converged in 2025. Every candidate started using AI to draft applications, polish resumes, and answer interview questions. And the EU AI Act started counting down to August 2026, when high-risk hiring AI must be designed, tested, documented, and monitored to standard. SCOUT was built for both shifts from the first line, not retrofitted to either.

CORE FEATURES

JD Pressure-Testing

Before a single candidate sees a job description, SCOUT tests it for ATS readability, bias-coded language, accidental disqualification of qualified candidates, and unclear requirements. Hiring teams catch the problems that depress applicant quality before posting, not three weeks into a stalled search.

Resume Scoring with Contact Extraction

Resumes flow through a staged scoring pipeline evaluating technical fit, experience depth, growth trajectory, and skill currency. Email, phone, LinkedIn, and portfolio links are extracted with confidence scores, so recruiter outreach lands on a verified address the first try.

Mandatory Screening Questions

Role-specific screening questions configured per posting with structured response capture. Candidates who clear screening move automatically into the interview stage. Candidates who do not get a courteous, documented decline with the legal posture intact.

Content-Only AI Interviewing

The AI interviewer evaluates what candidates say, not how they look, how they sound, or how they emote. Four hiring principles, designed in from the first line: no face, no voice, no emote, content only. Bias vectors that other platforms create are simply not in the architecture.

Per-Turn AI-Use Detection

Every candidate response is evaluated in real time for AI-assistance signals. Interviewers see the risk score alongside the answer, so a polished-but-shallow response gets noticed and followed up with a probing question, rather than scored as authentic depth.

Continuous Adverse Impact Analytics

Four-fifths rule, demographic deltas, and longitudinal bias drift are computed continuously rather than as a quarterly audit exercise. The first time an organization sees adverse impact appear is not the day a complaint is filed.

Audit-Ready Compliance Documentation

Every high-risk hiring AI deployment now requires a current technical file proving the system was designed, tested, and monitored to standard. SCOUT produces and maintains that file automatically as part of the hiring workflow, not as a separate compliance project months after the fact.

Jurisdiction-Aware Candidate Notices

NYC Local Law 144, Illinois AI Video Interview Act, Colorado AI Act, EU AI Act, and GDPR notifications each carry different obligations to the candidate. SCOUT surfaces the right notice to the right candidate based on their location, automatically.

KEY USE CASES

High-Volume Recruiting Pipelines

Organizations running hundreds of candidates per posting where screening at volume is the binding constraint on velocity. SCOUT handles the funnel end-to-end without compromising fairness or compliance.

AI-Aware Interviewing

Roles where candidate AI assistance during interviews is a real concern. Engineering, knowledge work, customer-facing positions. SCOUT's per-turn detection flags AI-assisted responses for live follow-up rather than scoring polish as depth.

EU Market Hiring Programs

Companies hiring in the EU after August 2026 need defensible documentation from the first job posting. SCOUT generates the dossier as part of the regular hiring workflow.

Multi-Jurisdiction Talent Acquisition

Enterprises hiring across NYC, Illinois, Colorado, and the EU face overlapping but distinct disclosure and audit requirements. SCOUT handles the jurisdictional differences automatically by candidate location.

Legacy Vendor Replacement

Organizations replacing legacy hiring AI tools whose compliance posture is no longer sufficient under EU AI Act requirements or expanding state-level rules, and whose interviewing approach was not designed for AI-using candidates.

APPLICATION INTELLIGENCE

ATLAS

One catalog. Every downstream system.

The application intelligence platform that crawls any web application, catalogs its complete surface, audits its user experience, and produces the artifacts every downstream consumer needs. Playwright code for QA automation. Grounding packs for AI agents. Sitemaps for documentation. SQL and JSON for everything else. One catalog feeds them all.

BUILT FOR QA leaders modernizing test automation · Enterprise architects and integration leads · AI program owners deploying agents into operational workflows · DevOps and SRE teams managing multi-environment estates

WHY NOW

Application intelligence used to mean three separate projects. QA wrote one application map for tests. AI wrote another for agent grounding. Documentation wrote a third for change management. Each was partial, each was out of date, none agreed. ATLAS produces all three from one continuously updated crawl, so every downstream system describes the application the same way: as it actually is right now.

CORE FEATURES

Multi-Environment Workspace

Manage configurations for production, test, DR, sandbox, and custom environments side by side. Register entirely new sandbox environments for proof-of-concept work without disrupting production maps. Switch contexts without losing state.

Adaptive Sitemap Crawler

Discovers every page, route, and interactive element across any application surface. Engine selection handles both modern dynamic web applications and traditional enterprise platforms, from custom internal portals to commercial SaaS.

Sign-On and UX Auditor

Audits authentication flows, navigation patterns, and user experience friction. Surfaces broken paths, dead-end workflows, and accessibility violations before customers find them.

UI Cartography

Maps interactive elements with the metadata downstream systems need: selectors, labels, types, dependencies, role-scoped visibility, and visual anchors. The catalog is queryable, structured, and complete, ready for any consumer.

Grounding Drift Audit

Applications change constantly. ATLAS re-crawls on schedule, diffs against the last known baseline, and flags every change. The silent drift that breaks automations and AI agents becomes visible before it causes incidents.

Playwright Codegen

Generates production-ready Playwright test code directly from the catalog. QA teams stop writing locators by hand and start running tests against the same source of truth every other downstream system uses.

Trace Self-Heal

When automation fails because a selector drifted, ATLAS traces the failure, locates the new selector, and heals the test. Self-maintaining automation replaces perpetual maintenance backlog.

Universal Output and MCP

Compiles the catalog into Markdown for documentation and RAG, SQL for relational systems, JSON for structured ingestion, and exposes it through the Model Context Protocol so any AI tool queries it directly. One catalog, every downstream shape.

KEY USE CASES

QA Automation Modernization

Replacing brittle hand-authored Playwright suites with generated tests that update automatically as the application evolves. Stop maintaining selectors. Start shipping coverage.

Enterprise AI Agent Deployments

Grounding AI agents in tenant-specific reality so they reason against facts instead of hallucinations. ATLAS produces the grounding data agent frameworks consume directly.

Multi-Environment Configuration Audits

Validating that production, test, DR, sandbox, and custom environments remain aligned, with drift flagged for investigation rather than discovered after a failed cutover.

Cross-Application Cataloging

From custom internal portals to commercial enterprise platforms, ATLAS catalogs the entire application stack rather than limiting itself to one vendor, one technology, or one consumer.

Pre-Migration Application Discovery

Source-system cataloging as the foundation of any migration program, ensuring the destination accounts for every page, field, flow, and integration actually present in production.

AI GOVERNANCE

SENTINEL

From AI inventory to modernization plan.

The AI inventory, opportunity, and advisory platform for the enterprise whose AI footprint has outgrown the spreadsheet. SENTINEL discovers and catalogs every AI deployment, identifies the highest-value automation opportunities, models stay-versus-migrate-versus-modernize economics, maps the inventory to every compliance framework that applies, and produces the evidence regulators and auditors will eventually require. One inventory, every downstream conversation.

BUILT FOR Chief AI Officers planning the AI roadmap · CIOs and CTOs planning database and platform modernization · CFOs evaluating AI spend and vendor consolidation · AI Governance leads producing compliance evidence · Compliance and risk officers preparing for the EU AI Act and parallel regulations

WHY NOW

Most enterprises now run more AI than they can list, and that list is what every conversation downstream depends on. The compliance team needs it for the EU AI Act and parallel regulations. The CFO needs it for vendor consolidation and spend rationalization. The CTO needs it for modernization planning. The CAIO needs it to decide what to build next. SENTINEL builds the inventory once and lets every downstream conversation start from the same source of truth.

CORE FEATURES

Comprehensive AI Discovery

Most enterprises cannot enumerate every AI touching their environment. SENTINEL discovers and catalogs every agent, copilot, embedded assistant, integration, and shadow AI deployment across the enterprise application and data surface, producing the first complete inventory most organizations have ever had.

Continuous Inventory Maintenance

Most governance projects produce a snapshot and decay from there. SENTINEL maintains a continuous inventory after the initial diagnostic, so the AI footprint stays current as new tools enter the environment and existing tools are decommissioned. The list never goes stale.

Agentic Workflow Opportunity Scoring

The platform identifies high-value candidates for agentic automation across finance, HR, supply chain, and customer operations. Period close, AP matching, supplier onboarding, and requisition approval are scored by frequency, cycle time, complexity, and automation feasibility. The output is a prioritized investment plan, not a list of ideas.

Three-Way Architecture Economics

Stay, migrate, or modernize in place are all viable strategies with different cost and capability trajectories. SENTINEL produces the comparative total cost of ownership across all three options, including the modernize-in-place path that vendor sales motions often omit.

Modernization Readiness Diagnostic

Database and platform modernization unlocks new AI capabilities, but only when the environment is prepared to receive them. SENTINEL assesses readiness and produces the architecture and capacity plan that makes the transition real rather than aspirational.

Vendor and Spend Consolidation Analysis

The AI inventory makes vendor overlap, capability redundancy, and spend concentration visible. CFOs and procurement leaders use SENTINEL output to consolidate the AI vendor footprint, retire duplicate tools, and rationalize spend across the estate.

Multi-Framework Compliance Mapping

EU AI Act, SOC 2, ISO 42001, NIST AI RMF, GDPR, CCPA, and NYC LL 144 each carry distinct evidence requirements. SENTINEL maps the AI inventory to each framework's obligations in a single pass, eliminating the duplicate work of preparing for each audit separately.

Governance Posture Scoring

A single composite score combines inventory completeness, documentation quality, monitoring coverage, and risk exposure into a number leadership can track quarter over quarter. The improvement trajectory becomes a board-level metric rather than a CAIO talking point.

COMPLIANCE FRAMEWORKS COVERED

EU AI Act

SOC 2

ISO 42001

NIST AI RMF

GDPR / CCPA

NYC LL 144

KEY USE CASES

AI Investment Planning

Identifying the highest-value automation opportunities across finance, HR, supply chain, and customer operations, then sequencing them by complexity, frequency, and expected return. The CAIO's roadmap, with evidence.

Vendor and Spend Consolidation

Producing the comparative AI spend, capability overlap, and risk exposure analysis CFOs need to consolidate the AI vendor footprint, eliminate duplicate tools, and rationalize total AI spend.

Modernization Decision Support

Modeling the workload uplift opportunity and the modernization economics across the full stay-versus-migrate-versus-modernize decision, so the architecture choice is made on evidence rather than vendor pressure.

Pre-Regulator AI Inventory

Producing the inventory the regulator will request before the request arrives. EU AI Act August 2026 deadlines, SEC AI disclosure rules, and state-level requirements are converging on the same need.

Shadow AI Discovery

Surfacing AI tools that entered the environment without IT approval through SaaS integrations, embedded copilots, or third-party data flows. The inventory the security team always wanted but never had.

KNOWLEDGE INTELLIGENCE

CODEX.

Read one document, or reason across forty-seven.

The document and corpus intelligence platform that translates, decomposes, summarizes, and reasons across whole bodies of knowledge. Available in cloud, on-premises, or sealed sovereign appliance modes, so the deployment matches the data classification.

BUILT FOR Legal and compliance teams reviewing high-volume document sets · Procurement teams analyzing supplier and partner contracts · Government and defense customers requiring on-premises or air-gapped deployment · International enterprises needing translation with structural fidelity

WHY NOW

Most document AI tools were built to read one document at a time. The decisions enterprises actually make span dozens or hundreds of documents simultaneously: every supplier contract, every regulatory filing, every audit response, the complete acquisition file, the entire policy library. CODEX was built to reason across the whole corpus in one pass, and to do it in cloud, on-premises, or sealed sovereign deployment so the architecture matches the data classification.

CORE FEATURES

Structurally-Aware Translation

Most translation tools flatten document structure, returning prose that lost its outline, tables, headers, and cross-references. CODEX preserves the structural fidelity of the source document while translating the content, so the output is usable without manual reconstruction.

Document Decomposition

Long, complex documents become navigable structural outlines that preserve hierarchy, section relationships, and dependencies. Legal teams, compliance officers, and procurement leads work from the structure of the document rather than from its prose.

Executive Summarization

Structured summaries with the key points, decisions, action items, and risks made explicit. The summary is not the prose compressed; it is the document reorganized for the reader who needs to act on it.

Cross-Document Reasoning

Most document AI tools work one document at a time, requiring users to assemble cross-document conclusions manually. CODEX reasons across an entire corpus simultaneously. Forty-seven supplier contracts, every audit report from the last seven years, the complete regulatory file. One question, one answer, drawn from everything at once.

Domain-Tuned Models

Generic models miss the language of specific industries and customer environments. CODEX fine-tunes on the customer's corpus, and the resulting model belongs to the customer rather than to the vendor. No retraining lock-in, no opaque data flows, no surprise cost escalation.

Cloud, Local, or Sovereign

Data classification varies across the enterprise. CODEX runs in cloud mode for general business documents, in local mode on customer infrastructure for sensitive corpora, and in sovereign appliance mode for classified or air-gapped environments. The deployment matches the data.

KEY USE CASES

Contract Portfolio Analysis

Procurement and legal teams analyzing supplier obligations, renewal exposure, and risk concentration across hundreds or thousands of contracts simultaneously rather than one at a time.

Regulatory Filing Synthesis

Cross-referencing every prior filing, regulatory response, and audit finding to ensure consistency in the current submission, with citations to source.

Acquisition Due Diligence

Reading the entire legal, financial, and contractual file of an acquisition target in days rather than weeks, with structured outputs the deal team can act on.

Public Records Compliance

Government entities responding to FOIA, public records, and discovery requests with documented exhaustiveness and defensible workflows.

Multi-Language Operational Documentation

International enterprises maintaining operational manuals, training materials, and compliance documents in many languages with structural fidelity preserved across translations.

DATA ARCHAEOLOGY

MOSAIC.

Read the fragments. Modernize the schema.

The legacy schema reconstruction and dead-vendor recovery platform for the systems no one understands anymore. MOSAIC recovers institutional knowledge from data evidence, produces a working specification, and hands the result off to modernization.

BUILT FOR Chief Information Officers facing forced migrations off dead-vendor databases · Banks, insurance carriers, and manufacturers running 20-plus year-old systems · Acquisition integration teams inheriting unfamiliar legacy databases · Public sector agencies modernizing decade-old core systems

WHY NOW

Every enterprise has at least one of these systems. A general ledger running on a database whose vendor disappeared a decade ago. A customer master file in a format only one retiring engineer can read. A claims system that survived three IT modernizations because nobody understood it well enough to replace it. The system runs. Moving it requires reading what no one remembers writing.

CORE FEATURES

Semantic Column Decoding

Legacy databases were built when storage was scarce and developers used eight-character column names. CUST_NM_1 and GL_AMT_USD_CY meant something specific to the original team, and that meaning was rarely documented. MOSAIC infers the semantic content of cryptic columns from data patterns, naming conventions, and structural relationships, recovering what the institutional knowledge contained.

Schema Reconstruction from Evidence

When the data dictionary is gone, the developers are gone, and the documentation never existed, MOSAIC reconstructs the schema from primary evidence: table contents, foreign key patterns, join behaviors, and statistical relationships between fields. The reconstructed schema becomes the working specification for whatever comes next.

Dead-Vendor System Recovery

Sybase, Informix, dBase, FoxPro, IDMS, Adabas, VSAM, IMS, and dead generations of Oracle and SQL Server still run critical operations at large enterprises. The vendor is gone or the version is unsupported. MOSAIC works directly on these systems, not against modern equivalents.

Data Lineage Reconstruction

Stored procedures, triggers, scheduled jobs, and batch processes form an undocumented graph of data movement through the system. MOSAIC traces the graph, identifying upstream and downstream dependencies that determine what can be migrated and in what order.

Continuous Migration Hand-Off

The reconstructed schema and lineage do not sit on a shelf. They feed directly into PRISM for the migration execution itself, making modernization one continuous engagement rather than two disconnected projects with a documentation gap between them.

KEY USE CASES

Forced Vendor Migration

When the vendor announces end-of-support and the customer has no internal expertise on the system, MOSAIC supplies the missing institutional knowledge so the migration can proceed on schedule.

Acquisition Integration

Acquiring companies inherit systems they did not build and did not document. MOSAIC accelerates the integration timeline by making the unfamiliar systems legible to the receiving IT organization.

Mainframe Modernization

IDMS, Adabas, VSAM, and IMS systems still run banking, insurance, and government operations. MOSAIC makes the modernization path tractable rather than speculative.

Data Warehouse Source Inclusion

Including legacy systems in modern data platforms requires understanding what the source data actually means. MOSAIC produces the understanding the modern platform needs.

Public Sector Modernization

State and local agencies running decades-old core systems where the original implementer is no longer available and the documentation is lost.

REGULATORY COMPLIANCE

OYEZ.

One regulator per quarter. One dossier per regulator.

The regulatory compliance platform for the chief compliance officer whose inbox now contains a new regulation every quarter. AI regulation, privacy law, cybersecurity rules, sectoral frameworks, ESG disclosure, trade compliance. OYEZ produces per-regulation dossiers built for what each regulator actually asks to see, in the language and structure each framework requires.

BUILT FOR Chief Compliance Officers across all regulatory domains · Data Protection Officers · Chief Information Security Officers preparing cyber attestations · General Counsel responsible for AI, privacy, and disclosure risk · Internal and external audit teams

WHY NOW

Compliance regimes are multiplying across every domain. EU AI Act in August 2026. CCPA amendments and state privacy laws. NYDFS cybersecurity rules and state CISO requirements. SEC AI disclosures and climate disclosure rules. CSRD reporting across Europe. HIPAA and GLBA evolution. Trade compliance and sanctions. NYC Local Law 144, Colorado AI Act, Texas DPSA. Each requires a different artifact. Each has a different deadline. Each lands in a different compliance officer's inbox.

CORE FEATURES

Per-Regulation Dossier Production

Generic governance reports are not what regulators ask for. OYEZ produces the specific artifact each regulator expects, in the language and structure that regulator's framework requires. Annex IV technical files for the EU AI Act, GDPR Article 35 DPIAs, NYC LL 144 audit reports, NYDFS cybersecurity attestations, SEC AI and climate disclosures, CSRD sustainability reports, HIPAA risk assessments, trade compliance attestations, and more, each produced to spec.

Regulatory Scope Detection

New regulations come into force constantly across every domain. Existing systems and data flows may fall into a new regulatory scope without the compliance team noticing. OYEZ monitors regulatory developments and flags when an existing system now requires new documentation.

Cross-Regulation Evidence Mapping

Evidence collected for one regulation often satisfies the requirements of another. The same SOC 2 controls support ISO 27001. The same DPIA inputs support GDPR and CCPA. OYEZ identifies the overlap, so the same evidence library serves multiple frameworks without duplicating effort.

Always-Current Audit Posture

Regulators increasingly ask for evidence with short response windows. OYEZ maintains dossiers in always-current state, so a regulatory request that arrives Tuesday morning can be answered Tuesday afternoon, not three weeks later.

Centralized Evidence Library

Policy documents, technical specifications, training data summaries, governance approvals, audit logs, security controls, and disclosure records live in one library indexed by regulation, jurisdiction, system, and date. The library is queryable, versioned, and tied to the specific dossiers it supports.

Deadline and Cadence Tracking

Effective dates, grace periods, annual reporting requirements, and continuous monitoring obligations each have their own cadence. OYEZ produces one calendar across every regulation the organization is subject to, regardless of domain.

KEY USE CASES

AI Compliance Programs

EU AI Act August 2026 obligations, NYC LL 144 expansion, Colorado AI Act, Illinois AI Video Interview Act. One library, one workflow, every required AI dossier.

Privacy Compliance Across Jurisdictions

GDPR, CCPA, state privacy laws, and sectoral privacy rules. The same evidence base serves DPIAs, transparency reports, and consumer requests across every applicable framework.

Cybersecurity Attestations

NYDFS Part 500, state CISO requirements, sector-specific cyber rules, SOC 2 evidence packages, and incident reporting obligations. The CISO and the compliance officer work from the same evidence.

ESG and Climate Disclosure

SEC climate disclosure rules, CSRD reporting across Europe, state ESG requirements, and supply chain sustainability attestations. The reporting work compounds; OYEZ scales with it.

Multi-Domain Regulator Response

When a state Attorney General, federal regulator, or sectoral authority opens an inquiry that crosses AI, privacy, cybersecurity, or disclosure boundaries, OYEZ assembles the cross-domain evidence package within the requested timeframe.

ITERIA'S FLAGSHIP PLATFORM

PRISM.

The complete modernization platform.

Iteria's flagship data modernization platform. PRISM connects, assesses, maps, archaeologically decomposes, classifies, and migrates the entire data estate from one engagement record, then uplifts workloads to native AI capabilities as part of the same motion.

BUILT FOR CIOs and CTOs running database modernization programs · Oracle Database administrators evaluating 23ai or 26ai migration paths · Enterprises with multi-cloud database estates · Systems integrators delivering large-scale modernization engagements

WHY NOW

Database modernization is no longer a lift-and-shift question. Modern target platforms make native vector search, AI-assisted query, JSON-Relational Duality, and Property Graph core database capabilities. The strategic decision is no longer whether to modernize but which workloads to uplift to which capabilities, and on what timeline. PRISM answers that across multi-cloud sources and targets.

CORE FEATURES

Multi-Source, Multi-Cloud Connectivity

PRISM connects natively to every major database family. Oracle (back to the earliest production versions), SQL Server, PostgreSQL, MySQL, DB2, Sybase, Informix, and more. The heterogeneous estate is treated as one engagement rather than a sequence of point migrations.

Workload Discovery and Risk Scoring

Before any migration runs, PRISM produces a complete workload analysis: dependency graphs, complexity scoring, risk ranking, and the prioritized order in which workloads should move. Every migration plan is built on evidence, not on assumption.

Native Eight-Phase Pipeline

Connect, Assess, Architecture Map, Package Archaeology, External Map, Patch Compiler, Taxonomy, and Portal. All eight phases live in one product, share one data model, and feed one engagement record. Phases do not hand off between tools; they progress within the same record.

Capability Uplift Scoring

Modern target platforms unlock new capabilities, but those capabilities only deliver value when the right workloads are matched to them. PRISM scores every workload against the destination's native capabilities and produces the prioritized uplift recommendations.

Lift-and-Uplift Automation

Migration is rarely pure lift-and-shift. PRISM identifies patterns during migration that should be modernized in place: fuzzy match logic becomes vector search, materialized JSON views become native duality views, complex multi-table joins become graph queries. Modernization happens during migration, not as a separate project afterward.

Engagement Command Center

One portal, one engagement record, one source of truth for the entire program. Status, risks, decisions, and artifacts live in one place from Phase 1 connect through Phase 8 cut-over.

KEY USE CASES

Multi-Cloud Database Consolidation

Enterprises with Oracle, SQL Server, and PostgreSQL workloads consolidating onto a single modern platform with a unified migration program rather than three separate projects.

Modern Platform Migration

Existing database customers planning a migration to current platform versions, including the workload uplift to native AI capabilities that justifies the migration economically.

End-of-Support Database Replacement

Forced migrations from databases approaching end-of-life or end-of-support, where the modernization timeline is driven by vendor sunset rather than internal priority.

Data Center Exit Programs

Multi-year programs moving entire database estates from on-premises infrastructure to cloud platforms, with PRISM as the single tool of record across the program.

Acquisition Database Integration

Combining acquired company database estates into the parent organization's modernization platform under one engagement record.