OrbisFramework

ENTERPRISE AI WORKFLOW INFRASTRUCTURE

The Infrastructure Layer for Complex Enterprise AI Workflows.

Stop building the engine. OrbisFramework delivers the AI orchestration, data integration, security, and decision infrastructure that complex enterprise workflows demand, so development focuses on workflow value, not foundation work.

100+

AI Models. Any domain. Per-step engine selection.

2 wks

To production. One developer. Enterprise-grade workflow.

Zero

Infrastructure sprints. The foundation is already built.

CONFIDENTIAL

THE ENTERPRISE AI BUILD PROBLEM

Every Complex AI Workflow Starts With the Same Trap

Before a single workflow runs, organizations attempting enterprise AI must first build the same foundation: security, orchestration, data integration, audit compliance. That foundation consumes the engineering capacity that should be building workflow value.

⏱️

The Infrastructure Tax

Authentication, authorization, RBAC, audit logging, multi-model AI management, and input security are all required before a single workflow step executes. This is typically 6 to 12 months of foundation work that produces nothing the business sees. Engineers build what must exist, not what creates value.

🔌

The Integration Wall

AI operating without live organizational data produces generic, low-value outputs. Connecting enterprise databases, knowledge systems, and external APIs securely is a separate engineering effort on top of the AI build. Each data source requires its own integration, credential management, and query design.

🔄

The Iteration Bottleneck

When AI prompt logic, model selection, and workflow sequencing live in code, every behavioral change requires a development cycle: development, testing, deployment, validation. Business iteration moves at engineering speed, not at the speed insight demands. The cost of experimentation becomes prohibitive.

The result: engineering capacity consumed by foundation work organizations build once, maintain forever, and rarely differentiate on, while the actual workflow value waits.

THE COST

6 to 12 months of engineering before value delivery. A full team. A significant budget. And still no workflow running in production.

THE RISK

Security gaps and compliance exposure emerge when security, audit, and input protection are bolted on rather than built into the foundation from the start.

THE OPPORTUNITY COST

Every month spent on foundation work is a month competitors with better infrastructure are deploying, iterating, and widening the gap.

What the problem demands

Infrastructure that is pre-built

Security, orchestration, audit, and data integration delivered as a foundation, not a build requirement.

Workflow logic that is configurable

AI behavior, model selection, and step sequencing managed through configuration, not code deployments.

Development that stays focused

Developer effort concentrated on workflow screens and user experience, not on rebuilding shared infrastructure.

WHAT ORBISFRAMEWORK ELIMINATES

The Foundation Is Already Built.

Every capability below is pre-built, production-grade, and inherited the moment development begins. None of it needs to be designed, built, tested, or maintained.

AI Orchestration Engine

Complete prompt configuration, chaining, and execution pipeline. Prompts are defined in the database: system instructions, templates, output format, model parameters. AI outputs map directly to inputs of subsequent steps. Sequential, parallel, or conditional execution with automatic retry and fallback on model failure. No code required to change AI behavior.

Live Data Integration (RAG)

Multi-database connectivity across SQL Server, PostgreSQL, MySQL, and ODBC sources. Secure credential management via Azure Key Vault and AWS Secrets Manager. Parameterized queries inject live organizational data directly into AI prompt context at runtime. Multiple data sources can be combined per prompt step.

Multi-Model AI Management

Access to 100+ AI models through OpenRouter. Per-step engine selection assigns the optimal model to each workflow stage. Primary and fallback engine configuration, round-robin load balancing, and cost tracking per model. Private LLM endpoint support for air-gapped and data-sovereign deployments.

Enterprise Security

JWT authentication with HTTP-only cookie storage. RBAC with cascading privilege levels (read, update, deactivate, create). Permission aggregation across roles. Full input protection: XSS, SQL injection, prompt injection, command injection, and path traversal are all validated before any processing occurs. Brute force protection with progressive lockouts.

Audit and Compliance

Every action timestamped and attributed. Every AI execution logged with full inputs, outputs, model used, token counts, and user identity. Authentication events, session activity, and failed access attempts all recorded. A complete compliance trail that requires no additional instrumentation.

Structured Output Management

Define expected AI response schemas at field level: objects, arrays, nested structures, and primitive values. Display labels, visibility controls, sort and filter capabilities. Output fields map directly to subsequent prompt inputs, enabling complex multi-stage AI pipelines without additional code.

What this means in practice

Start at the workflow

No security sprint. No auth build. No data integration design. Development begins at the workflow itself, not the plumbing underneath it.

Iterate without deploying

AI behavior, model selection, prompt parameters, and step sequencing are all configurable through the database. No deployment cycle required to refine the workflow.

Inherit enterprise grade

Every workflow built on OrbisFramework inherits production-grade security, audit logging, and multi-model AI management from day one.

WHAT GETS BUILT

Workflow Logic. Screens. Decisions.

With the infrastructure inherited, development effort concentrates on three areas: configuring the workflow logic, building the user experience, and defining how AI-scored decisions gate progress through the workflow.

WORKSTREAM 01

Workflow Configuration

Defined entirely through database configuration. No code deployments required to iterate on AI behavior, step sequencing, or data sources.

- Define processes, categories, and steps
- Attach and tune AI prompts: model, temperature, tokens, format, output schema
- Connect RAG data sources to specific steps and map query parameters
- Map outputs from one AI step to inputs of the next
- Reorder, modify, or extend workflow steps without a deployment cycle
- A/B test different prompt configurations against the same step

WORKSTREAM 02

Screen Development

Built against established framework coding standards and patterns. AI coding tools accelerate development dramatically, reducing weeks of effort to days.

- Screens built to framework standards and component patterns
- AI coding tools (Claude Code, Copilot, and others) accelerate development
- Developer focus stays on workflow UX and logic, not infrastructure
- Consistent, maintainable codebase across all workflow screens
- Backend services, auth, permissions, and audit are already wired
- Screens connect to a fully operational API from day one

WORKSTREAM 03

Scoring and Decision Gates

AI-augmented workflows do not just generate outputs. They score them, gate progression, and put the human in the loop at the right moment.

- AI scores outputs against defined quality thresholds
- Below threshold: workflow loops, refines, and iterates automatically
- Above threshold: workflow advances to the next stage
- Human review triggered at defined decision points
- No AI output advances without human approval where it matters
- Scoring history tracked and attributed for audit purposes

HOW SCORING AND DECISION GATES WORK

AI Executes
Output Generated
AI completes the workflow step and produces a structured, schema-validated output

→

AI Scores
Quality Assessment
Output scored against threshold. Gaps and weaknesses are identified with specificity.

→

Gate Decision
Threshold Check
Below threshold: loop back, refine, re-score. Above threshold: advance to next stage.

→

Human in Loop
Human Approval
At defined points, human reviews, revises, and approves before the workflow continues.

AI augments every stage. Humans remain in control at every decision point.

AI IN THE FRAMEWORK

What AI Actually Does at Each Stage

Beyond model selection and API calls: here is what AI delivers inside an OrbisFramework workflow, expressed as capabilities, not infrastructure components.

AI That Thinks in Stages

Complex reasoning pipelines where each AI step builds on the structured output of the last. No glue code, no manual data passing. The framework routes outputs to inputs automatically across as many stages as the workflow demands. Each stage is a focused, purposeful reasoning task.

AI That Knows Your Data

Live database context injected into every prompt at runtime. AI responses are grounded in actual organizational records: customer data, product catalogs, knowledge bases, transaction history. Not generic training data alone. The AI always has the context it needs to produce relevant, specific outputs.

AI That Challenges Before It Builds

Socratic planning loops surface hidden assumptions, expose gaps, and pressure-test alignment before the workflow commits to a direction. AI asks the questions that reveal whether the foundation is sound before execution invests in building on top of it.

AI That Critiques Its Own Work

Built-in critique cycles evaluate AI outputs against defined quality standards: clarity, completeness, rigor, logical consistency, or domain-specific criteria. Critique findings feed directly into enhancement steps, producing progressively stronger outputs without manual review overhead.

The Right Model for Each Task

Per-step engine selection means a workflow is not constrained to one model for everything. Reasoning-intensive steps use reasoning-optimized models. Generation steps use generation-optimized models. Each stage gets the most effective and cost-efficient model for what it needs to do.

AI Within Your Security Perimeter

Private LLM endpoint support enables connection to self-hosted models (Ollama, vLLM, LocalAI), private cloud deployments (Azure OpenAI, AWS Bedrock), or air-gapped infrastructure where no data crosses a public network boundary. Full AI capability with zero external data exposure.

WHAT IT WAS BUILT ON

The Hardest Workflow First.

OrbisFramework was not designed in the abstract. We took the most complex AI-augmented workflow we could conceive: academic research, from initial idea to publication-ready manuscript, and built it completely, with full human-in-the-loop controls at every stage. Every infrastructure requirement that emerged from that process became part of the framework. What survived became the foundation everything else is built on.

THE FOUNDATION

AI-Augmented Research Development Platform

The workflow that built the framework

- End-to-end research lifecycle from initial idea to publication-ready manuscript
- 32 steps across 5 phases with 8 Socratic planning processes
- Fully automated literature search pipeline: database query, metadata enrichment, dual-AI screening, deep review
- Theory suggestion and verification via OpenAlex API integration
- 20-section manuscript development with per-section plan, outline, draft, critique, enhance workflow
- Worldwide multi-user collaboration with role-based access control

AUTOMOTIVE DIAGNOSTICS

AI-Guided Vehicle Diagnostic and Repair Workflow

Built in 2 weeks

- Multi-step diagnostic pipeline from fault code capture through repair order generation
- RAG integration pulling live OEM technical bulletins, parts databases, and vehicle records
- AI-guided technician decision support with Socratic refinement before repair commitment
- Structured diagnostic output schemas for reports, repair orders, and compliance documentation
- Scoring gates that validate diagnostic confidence before recommending repair actions
- Role-based access across technicians, service advisors, and fleet managers

EDUCATION

AI-Augmented Learning Management System

Built in 3 weeks

- Full course lifecycle from enrollment through assessment, grading, and certificate generation with SPARK pedagogy integration (Strategic Challenges, Peer-to-Peer Review, AI Augmented Learning, Relationship Building, Kinetic Continuous Learning)
- AI-powered essay grading with structured rubrics, percentage-based scoring thresholds, and detailed feedback generation per assessment item
- Hierarchical course navigation across courses, modules, steps, and content pages with Vimeo-hosted video, rich text lessons, and embedded assessments
- Admin course builder with enrollment management, student progress tracking, and comprehensive reports and analytics dashboards
- Dual revenue model supporting tiered platform subscriptions and individual course purchases with full payment lifecycle management
- Same framework infrastructure. Zero additional foundation work required.

Three different domains. Three different teams and industries. The same infrastructure underneath. Because the framework was forged on the hardest workflow first, everything after it is faster.

Same infrastructure. Any workflow. Any domain.

ARCHITECTURE AND SECURITY

Production-Grade. Enterprise-Ready.

Technical Stack

FRONTEND

React 19.1 · Vite 4.5 · TailwindCSS 4.1
TanStack Query 5.85 · React Router 7.8 · Axios 1.11

BACKEND

Node.js · Express 5.1 · SQL Server (MSSQL 11)
Azure Key Vault · Azure Service Bus · bcrypt · Helmet

AI AND INTEGRATION

OpenRouter (100+ models) · Private LLM endpoints
Ollama · vLLM · Azure OpenAI · AWS Bedrock
OpenAlex API · CrossRef · Custom endpoints

DATA SOURCES (RAG)

SQL Server · PostgreSQL · MySQL · ODBC-compatible
AWS Secrets Manager · Azure Key Vault · SSL/TLS enforced

DEPLOYMENT OPTIONS

SaaS · Managed Private Cloud · Azure / AWS
Air-Gapped · Private LLM

Air-gapped: Runs entirely within the organization's physical or network boundary with no public internet connection. All AI processing occurs on internally hosted models. Data never leaves the security perimeter. Required for defense, intelligence, healthcare, and regulated financial environments.

Security Architecture

AUTHENTICATION

JWT-based with HTTP-only cookie storage, eliminating client-side token exposure. Session management with configurable inactivity timeouts. Token refresh with full audit logging. Brute force protection with progressive account lockouts that escalate with repeated failures.

AUTHORIZATION · RBAC

Table-level permissions with four cascading privilege levels: read, update, deactivate, create. Permission aggregation across multiple roles ensures users inherit the highest applicable permission. Project-specific access enforcement applied throughout the entire request lifecycle.

INPUT PROTECTION

All inputs validated before any processing occurs. XSS payload detection, SQL injection prevention, prompt injection blocking, command injection blocking, and path traversal prevention are all enforced at the request pipeline level, not the application layer.

AUDIT AND COMPLIANCE

Every action timestamped and attributed to a named user. Every AI call logged with the model used, full inputs, outputs, token counts, and cost. Authentication metrics, session activity, and all failed access attempts recorded. A complete compliance trail requiring no additional instrumentation or tooling.

WHY ORBISFRAMEWORK

Stop building the engine. Build the workflow. Achieve the benefits.

OrbisFramework is the infrastructure layer that complex enterprise AI workflows require. Built by taking the hardest workflow imaginable and solving it completely, it is a foundation no organization should have to build twice.

ELIMINATE

The Foundation Sprint

Security, RBAC, AI orchestration, audit, multi-model management: pre-built. 6 to 12 months of foundation work gone before day one. Engineering starts at the workflow, not the plumbing.

DEPLOY

With Velocity

One developer. An automotive diagnostic workflow in 2 weeks. AI coding tools accelerate screen development further still. Workflows reach production in weeks, not quarters.

OWN

The Workflow

Configure the logic. Build the experience. Run on enterprise-grade infrastructure your organization controls: SaaS, private cloud, or fully air-gapped within your security perimeter.

The benefits organizations achieve

- Complex, multi-step AI workflows in production within weeks, not years
- AI grounded in live organizational data, producing specific and relevant outputs
- Workflow behavior that iterates through configuration, not code deployments
- Enterprise-grade security and compliance inherited from day one
- Human judgment at every decision point, supported by AI scoring and gates
- Any domain. Any workflow complexity. The same foundation underneath.

Any domain. Any workflow complexity.
The infrastructure does not change. Only the workflow does.

