VERIAX
ENTERPRISE AI ORCHESTRATION PLATFORM

Autonomous AI agents for end-to-end business processes

Developed by THUX CODE - 30 years of enterprise software

ISO 9001:2015 · ISO 27001:2022 · Deployment in 12 weeks · Three perimeter-agnostic modes

VERIAX
ENTERPRISE AI ORCHESTRATION PLATFORM

Autonomous AI agents for end-to-end business processes

VERIAX orchestrates complete workflows through autonomous AI agents within the perimeter defined by the customer. Three deployment modes — on-premise, private cloud (customer's AWS, Azure, GCP), hosted by THUX CODE — for perimeter-agnostic data sovereignty. Model-agnostic architecture for independence from LLM providers. Native governance for AI Act, GDPR and NIS2 compliance.

The pilot-to-production gap

Three structural bottlenecks keep most enterprise AI projects stuck in pilot.

DATA PERIMETER	INTERNAL SKILLS	OPERATIONAL GOVERNANCE
Many organizations process sensitive data that cannot leave their corporate infrastructure. Mainstream SaaS solutions require a public cloud connection. In-perimeter deployment removes the constraint but introduces infrastructure complexity that calls for specialized skills.	Enterprise AI solutions require LLM orchestration, prompt engineering, vector databases and infrastructure management. Most organizations don't have these skills in-house and cannot afford 6 to 12 month hiring cycles.	The AI Act requires a complete audit trail, configurable human oversight and lifecycle documentation of the AI system. GDPR and NIS2 add sector-specific requirements. Compliance cannot be an add-on: it must be native to the architecture.

Built by a team that has been shipping enterprise software for 30 years

30 YEARS	MATURE ENGINEERING	ISO 9001 ISO 27001
Enterprise heritage	From semantic search to agents	Quality · Information security
THUX CODE has been building solutions for Finance, Manufacturing and Public Administration since 1996. VERIAX is not a pivot: it's the natural evolution of a mature engineering framework applied to enterprise LLMs.	From the 2019 semantic search prototype to emotion recognition projects, conversational RAG and document management for enterprise customers, each iteration has built on the one before. When LLMs became reliable for critical contexts, infrastructure and expertise were already in place.	Verifiable certifications in quality management and information security. VERIAX runs the quality management system and software development processes of THUX CODE internally: the framework is our first user.

Three deployment modes, one data sovereignty principle

VERIAX
ENTERPRISE AI ORCHESTRATION PLATFORM

VERIAX is perimeter-agnostic: execution stays inside the perimeter defined by the customer in all three modes. It's not a contractual promise, it's the architectural structure of the framework.

ON-PREMISE	PRIVATE / PUBLIC CLOUD	HOSTED BY THUX CODE
VERIAX runs on the customer's infrastructure (virtual machines, bare metal, Docker containers). Data never leaves the corporate data center.	Deployment on the customer's private cloud or Public cloud (AWS, Azure, GCP). Compute resources stay under the customer's IT control, with governance and billing on the corporate cloud account.	Hosting on infrastructure managed by THUX CODE for those who don't want to run the deployment themselves. The same data sovereignty constraint is preserved contractually, with EU data residency and full traceability.

How it works

VERIAX follows an input → orchestration → output model.

On input, the framework connects to data sources already in place within the organization: relational databases, document repositories (PDF, Word, Excel, PowerPoint), business systems, email and corporate collaboration platforms, with no data migration involved.

Orchestration is handled by a system of agents configurable across four autonomy levels: from simple agents for direct queries to agents capable of multi-step planning and self-validation. Advanced RAG (semantic search plus keyword) and Graph RAG anchor responses to internal sources and enable multi-hop reasoning over complex relationships. Every agent operates within defined limits, with mandatory human approval (HITL) on irreversible operations.

On output, workflows produce documents, reports, data analyses, database queries and notifications without requiring technical skills from end users. Every action is logged in an immutable, tamper-evident audit trail, a direct requirement of the AI Act on decision auditability.

Key benefits

From POC to production in 12 weeks. Most enterprise AI projects require 12 or more months of in-house development. VERIAX includes a structured path: Proof of Value in 4 to 6 weeks on the first real use case, then scaling to production within 12 weeks.

Perimeter-agnostic data sovereignty. No corporate data leaves the defined perimeter. VERIAX deploys in three modes: on-premise, customer private cloud or Public Cloud, hosted by THUX CODE, with zero transmissions to third-party servers.

AI Act-ready compliance. Append-only and tamper-evident audit trail (AI Act Art. 13), documented HITL (Art. 14), AI monitoring (Art. 17) and documentation templates (Art. 53) are part of the baseline architecture. Native PII anonymization. Final compliance depends on the individual use case and requires legal review.

Measurable ROI from the first use case. Finance reference scenario: a team of 10 analysts spending 20 hours/month on document production. Post-VERIAX: 25 hours/month and 45 minutes per report (validation of agent output). Net saving: 175 hours/month, with a recovered value estimated at €94,500/year. Transparent cost model, no per-user fees.

Native governance for every user profile. Granular RBAC, configurable rate limiting per agent and per-department access control. Administrators manage the framework without touching code, while domain managers configure agents and workflows through a no-code interface.

Integration with the existing IT stack. Hundreds of native connectors for databases, CRM, ERP and document systems; integrable with virtually any system. Model-agnostic architecture: change the LLM at runtime via the UI, with no service interruption and no recoding. If an LLM provider becomes unavailable, automatic fallback ensures operational continuity.

The enterprise buyer committee

In regulated sectors, the VERIAX adoption decision involves a buyer committee with at least five roles. For each, the dominant pain and the corresponding value.

PERSONA	PRIMARY PAIN	VERIAX VALUE
CTO / IT Director	Too many AI POCs never reach production due to missing governance, integration and a structured path.	12-week time-to-production, native integration with existing systems, a complete framework rather than a kit to assemble.
COO / Operations Director	Slow document and decision processes, reliance on qualified staff for low-value repetitive activities.	End-to-end agentic workflows, domain managers operating autonomously through no-code, measurable ROI from the first use case.
CFO / Finance	Every technology investment must show measurable ROI within 12 months; AI pricing is unpredictable.	ROI estimate verifiable from Proof of Value (4 to 6 weeks), transparent cost model, quantified savings on document processes.
CISO / Security	No sensitive data can be processed outside the controlled perimeter; security audits required.	Three perimeter-agnostic deployment modes, tamper-evident audit log, granular RBAC, HITL on irreversible operations, ISO 27001:2022 certified vendor.
Compliance Officer / DPO	AI adoption must be demonstrably compliant with the AI Act, GDPR and sector regulations.	Complete audit trail (Art. 13), documented HITL (Art. 14), native PII anonymization, documentation templates (Art. 53), dedicated GDPR vertical skill.

VERIAX
ENTERPRISE AI ORCHESTRATION PLATFORM

Sector-distinctive use cases

Each use case shows end-to-end orchestration leveraging capabilities that set VERIAX apart from the alternatives the market evaluates: multi-hop reasoning over knowledge graphs (Graph RAG), threshold-configurable human oversight (HITL), automatic guardrails for sensitive data, different LLMs for different tasks with native fallback, multi-channel integrations, supervisor agents that coordinate specialized sub-agents.

CONSULTING & PROFESSIONAL SERVICES · **Cross-project knowledge management with Graph RAG**

Consulting firms accumulate knowledge spread across projects, clients, deliverables and presentations. The VERIAX agent builds a knowledge graph of project-client-skill-deliverable relationships, identifies multi-hop patterns on analogous requests, recommends the optimal team based on skills demonstrated in similar contexts, and generates the initial brief drawing from past deliverables with precise citations to source documents.

Outcome. From client request to operational brief, allocated team and deliverable draft in 2 hours instead of 3 days of manual research. Corporate knowledge that's actually reusable.

Distinctive capabilities · **Graph RAG** · **Self-query metadata** · **Semantic cache** · **Precise citations**

RETAIL & E-COMMERCE · **Multi-channel sentiment analysis with automated escalation**

Complaints and returns come in via email, chat, web forms and social media with heterogeneous formats. The agent ingests the flows through multi-channel connectors, classifies severity and sentiment with configurable guardrails (critical keywords such as "lawyer" or "compensation" trigger immediate escalation), for critical cases it notifies the manager and opens a priority ticket, for standard cases it uses a semantic cache with responses pre-approved by the legal team.

Outcome. Average response time: 15 minutes vs. 4 hours. Critical escalation: immediate vs. 1 to 2 days. Recurring problem patterns identified in real time, not at quarter end.

Distinctive capabilities · **Multi-channel connectors** · **Escalation guardrails** · **HITL** · **Semantic cache**

MANUFACTURING & INDUSTRY · **Non-conformity root cause analysis with Graph RAG over 15 years of production data**

A quality non-conformity calls for analysis of data spread across historical NC records, machine parameters, procedures and audit reports. The agent rebuilds the context drawing from the real-time production database, builds a relational knowledge graph over 15 years of NCs (product → process → cause → corrective action → verified effectiveness), identifies analogous NCs through semantic search plus self-query on metadata (year, production line, supplier) and proposes the corrective action based on verified historical evidence, with audit-ready ISO 9001 documentation generation.

Outcome. Average NC analysis time: 1.5 days vs. 4 days. Corrective actions grounded in historical evidence rather than intuition. Audit documentation always complete and traceable.

Distinctive capabilities · **Graph RAG** · **Real-time DB** · **Self-query metadata** · **OCR** · **ISO 9001 PDF generation**

VERIAX
ENTERPRISE AI ORCHESTRATION PLATFORM

FINANCE & INSURANCE · **Claims underwriting with anti-fraud guardrails and threshold-based HITL**

A car insurance claim requires policy coverage verification, documentation analysis (damage photos, expert report, repair quotes), fraud pattern detection, settlement calculation, multiple threshold-based approvals. The agent verifies coverage via SQL query on the policy database, analyzes damage photos with vision plus OCR on quotes, compares with historical fraud patterns via Graph RAG, applies automatic guardrails (amounts above a threshold or suspicious patterns trigger automatic blocking of sensitive data sent to the LLM plus mandatory HITL) and routes cases to the correct approver.

Outcome. Below-threshold settlements: automatic in 2 hours vs. 3 days. Fraud detection: 40% more suspicious patterns identified. Complete audit trail for regulatory compliance. Customer-sensitive data never exposed to cloud LLMs.

Distinctive capabilities · **Vision + OCR** · **Graph RAG** · **Sensitive data guardrails** · **Threshold HITL** · **Audit trail**

ENERGY & UTILITIES · **Field technician dispatch with route optimization and multi-system coordination**

An electricity service outage requires locating the fault on the network, assigning the available technician with the right skills, optimizing the route considering other interventions in the area, communicating the ETA to all affected customers and updating status across NOC, CRM and field service systems. A supervisor agent coordinates specialized sub-agents: technician scheduling with skills matching, route optimization against real-time traffic, customer communication via mass SMS, NOC ticket opening.

Outcome. Dispatch time: 8 minutes vs. 35 minutes. Optimized routes with -15% intervention time. Proactive communication: 100% of affected users informed within 5 minutes. Automatic SLA reports for the regulator.

Distinctive capabilities · **Supervisor agent** · **Network topology Graph RAG** · **Real-time DB** · **Custom tools**

EDUCATION & RESEARCH · **Course accreditation with multi-source regulatory verification**

Accrediting a new course requires verifying compliance with national legislation (Ministry PDFs), regional guidelines (agency websites), internal quality standards and scientific committee resolutions. The agent builds a regulatory knowledge base through multi-source ingestion, applies self-query on metadata to filter by year in force, region and course type, verifies compliance point by point with advanced RAG and precise citations to regulatory references, and generates a compliance checklist plus gap analysis plus resolution draft in an institutional template.

Outcome. Accreditation time: 3 days vs. 3 weeks. Zero risk of regulatory non-compliance: every requirement cited with a verifiable source. Complete audit trail for certifying bodies.

Distinctive capabilities · **Multi-source ingestion** · **Self-query metadata** · **Precise citations** · **Committee approval HITL**

PUBLIC SECTOR & GOVERNMENT · **Case processing with GDPR privacy guardrails and AI Act transparency**

A building permit case requires verifying compliance with the municipal urban plan, building regulations and landscape constraints, opinions from technical offices, generating a compliant decision, notifying the citizen. Sensitive data must remain on-premise for GDPR. The agent applies automatic guardrails (masking of national identification numbers, addresses, personal data before processing), verifies compliance via RAG on the urban plan and regulations, a coordinator supervisor routes opinions to the competent offices, and the

output is a decision in a municipal template with precise regulatory citations and an AI Act compliant audit trail.

Outcome. Case processing time: 15 days vs. 45 days — procedural deadlines finally met. Zero data breaches on citizen data thanks to automatic guardrails. Full AI Act transparency: every decision traceable and motivated.

Distinctive capabilities · **Privacy guardrails** · **Supervisor agent** · **Multi-source RAG** · **AI Act audit trail**

TRANSPORT & LOGISTICS · **Multi-country customs documentation with provider fallback**

An international shipment outside the EU requires customs documentation compliant for transit and destination countries (commercial invoice, packing list, certificates of origin, value declarations). Regulations change by country; documentation errors trigger customs blocks and penalties. The system must always be operational. The agent retrieves up-to-date regulations via web search on official customs sites, generates documentation with operational skills specialized by country (a US invoice differs from a China invoice), uses different LLMs for different tasks (lightweight model for goods classification, advanced model for critical regulatory compliance verification), and in case of provider unavailability triggers automatic fallback to the alternative provider.

Outcome. Documentation errors: -92% (from 8% to below 1%). System downtime: zero. Penalty costs for customs blocks: -€180,000/year. Above-threshold shipments: human oversight guaranteed with audit trail.

Distinctive capabilities · **Multi-provider with fallback** · **Country-specific skills** · **Regulatory web search** · **Threshold HITL**

HEALTHCARE & PHARMA · **Clinical pathway with video analysis and healthcare data guardrails**

An oncology patient follows a multi-step pathway across visits, diagnostic exams (imaging), reporting, multi-disciplinary consultations, therapy, follow-up, with data spread across the electronic health record, PACS, PDF reports and Word minutes. Healthcare data is highly sensitive (GDPR plus regional healthcare regulations). The agent rebuilds the clinical timeline via Graph RAG on visit-exam-report-therapy relationships, analyzes imaging with OCR on reports plus video analysis on CT/MRI scans, applies healthcare guardrails for masking of identifying data and on-premise imaging processing with a self-hosted open source model; every therapeutic decision requires the oncologist's approval with a complete audit trail.

Outcome. Tumor board preparation time: 45 minutes vs. 4 hours. Informational continuity: 100% of relevant clinical data consolidated. GDPR compliance guaranteed (data never on public cloud). Medical responsibility preserved: mandatory HITL on every clinical decision.

Distinctive capabilities · **Clinical pathway Graph RAG** · **Vision / video analysis** · **Healthcare guardrails** · **Self-hosted open source model** · **Mandatory HITL**

LEGAL & COMPLIANCE · **GDPR compliance maintenance with event-driven workflows and dual oversight**

Companies operate with a privacy system set up once at inception and then essentially static, while the Record of Processing Activities, appointments, privacy notices and DPAs should update at the pace of business life. The agent triggers automatically on real business events (new client, supplier, employee, software adoption, data subject request, suspected data breach) qualifies the relationship under GDPR, updates the RoPA in line with Article 30, generates internal appointments, prepares or verifies the DPA, flags privacy notice updates. Two oversight agents work independently above the operational layer: one watches internal

consistency of the privacy system (controller-processor chain, appointments vs. org chart, DPA coverage), the other monitors external regulatory sources (Garante, EDPB, CJEU) and flags developments to the consultant for validation.

Outcome. From event to updated documentation and complete accountability evidence in minutes instead of waiting for the next consultant review. Compliance shifts from documentary promise to living proof.

Distinctive capabilities · **Event-driven workflows** · **Two-layer oversight** · **HITL on legal qualifications** · **EU data residency** · **Native audit trail**

Why VERIAX over the alternatives

- Unlike enterprise conversational AI assistants, VERIAX doesn't respond to point queries but executes complete end-to-end workflows without continuous supervision. In regulated sectors, conversation is a means; the operational result is the goal.
- Unlike cloud AI infrastructure platforms (building blocks that require custom development), VERIAX is a complete framework. Ready-to-use agent orchestration, native governance, vertical modules for documentation, data and workflows operational from day one, without requiring an internal AI team.
- Only VERIAX combines perimeter-agnostic data sovereignty (on-premise, private cloud, hosted), AI Act-ready governance and a 12-week time-to-production without tying the customer to a single LLM provider.
- Unlike custom solutions developed in-house, VERIAX is built on 30 years of enterprise experience at THUX CODE. Source code in escrow on request. Contractual SLAs on availability and response times.
- VERIAX doesn't require an internal AI team to operate. Setup managed by THUX CODE; day-to-day operations in the hands of an IT Admin and domain managers, without programming skills.

Why THUX CODE

THUX CODE is the AI division of THUX Srl, an Italian software house with 30 years of activity. VERIAX was developed in continuity with a track record of enterprise engineering for Finance, Manufacturing and Public Administration.

To enterprise customers, THUX CODE offers:

- Verifiable ISO 9001:2015 and ISO 27001:2022 certifications.
- A project team identified by name in the offer, before signing. The same people who analyze the context lead delivery and user training.
- Contractual SLAs on availability and response times, dedicated post go-live support.
- Source code escrow on request. Transparent product roadmap with scheduled releases.

VERIAX
ENTERPRISE AI ORCHESTRATION PLATFORM

Technical analysis of the use case

A structured session (2 to 4 hours) in which THUX CODE analyzes the operational context, identifies the use case with the fastest ROI and produces a technical roadmap for the first deployment.

- Mapping of priority operational pain points and processes that are candidates for orchestration.
- Identification of the quick-win use case for the first deployment.
- Preliminary ROI estimate on the selected use case.
- Technical roadmap: objectives, timeline, KPIs and technical dependencies.

Before the session, we run a structured AI Readiness Assessment with the client's team: the technical analysis starts from the real context, not from scratch. The THUX CODE team is named in the commercial offer before signing and stays in charge through to end-user training.