◆ VEROXOS

UNGATED · NO FORM · 2026

— AI GOVERNANCE PAPER

Ask. See. Act. Governed.

How Veroxos applies the NIST AI Risk Management Framework and prepares for the EU AI Act across the Veroxos AI Copilot — the conversational layer over telecom, mobility and asset-lifecycle data.

- → Audit-log architecture: what is recorded when the AI recommends or acts
- → No-training-by-default and the customer opt-in workflow
- → Human-in-the-loop approval design for every consequential action
- → NIST AI RMF function mapping and EU AI Act readiness position

1 · WHY THIS PAPER EXISTS

Enterprise buyers should not have to take AI governance on faith.

The Veroxos AI Copilot is a conversational layer over an enterprise's full Veroxos data model: invoices, contracts, inventory, usage, tickets and disposition records across telecom, mobility and asset lifecycle. It answers plain-English questions ("which carriers over-billed us this quarter, and by how much?"), surfaces anomalies before they cost money, and drafts actions — a carrier dispute, an optimisation change, a report — for a human to approve.

That capability raises exactly the questions a procurement, security or legal team should ask: What data does the AI see? Is our data used to train models? What is logged when the AI recommends or takes an action? Who approves what? How does this map to recognised frameworks? This paper answers those questions specifically, so the answers can be evaluated — and contractually referenced — rather than taken on faith.

Scope

- → **In scope:** the Veroxos AI Copilot — conversational query, anomaly detection and explanation, recommendation generation, and drafted actions across TEM, Mobility Management and Reverse Logistics data.
- → **In scope:** the governance controls around it — audit logging, data handling, model-training policy, human approval workflow, and access control.
- → **Out of scope:** general platform security controls not specific to AI (covered by the Veroxos Security & Compliance overview and the ISO 27001 ISMS), and any customer-side obligations that depend on how a customer deploys outputs downstream.

Our starting position

Veroxos is operated under StableLogic's ISO 27001-certified information security management system. A SOC 2 Type 2 audit is in progress (expected FY28). The platform is designed HIPAA-ready and GDPR-compliant, applies NIST SP 800-88 to media sanitisation in reverse logistics, and extends the same governance discipline to its AI layer: the controls in this paper are managed as part of the ISMS, not as a separate, informal effort.

A note on honesty in AI claims: nothing in this paper claims certification against the NIST AI RMF (a voluntary framework with no certification scheme) or "EU AI Act compliance" as a finished state. We describe what is implemented, what is by design, and what is in progress — and we label which is which.

2 · GOVERNANCE PRINCIPLES

Five principles the Copilot is built around.

1. The AI drafts; a human decides.

No consequential action — a carrier dispute, a plan change, a disconnection request, an external communication — is executed by the Copilot autonomously. The Copilot prepares the action with its evidence attached; a named, authorised person approves or rejects it. Read-only answers (queries, reports, explanations) do not require approval, and are logged.

2. No training on customer data without explicit opt-in.

Customer invoice, contract, inventory and usage data is never used to train or fine-tune foundation models by default. Participation in model improvement is a deliberate, contract-level opt-in with a defined scope, and it can be withdrawn (Section 5).

3. Every AI interaction leaves an audit trail.

Questions asked, data scopes touched, answers given, anomalies flagged, recommendations made, approvals granted or refused — each is a logged event with actor, timestamp and evidence reference (Section 4). If it cannot be reconstructed later, it did not happen in a governed system.

4. The Copilot cites its evidence.

Answers reference the underlying records — the invoice lines, contract clauses and usage rows behind a figure — so an analyst can verify rather than trust. Where confidence is low or data is incomplete, the Copilot says so instead of guessing.

5. Least privilege applies to the AI, too.

The Copilot operates within the querying user’s existing role-based permissions. It cannot read data, or draft actions against data, that the signed-in user could not access through the ordinary interface. Cross-customer data access is architecturally excluded: each engagement’s data model is segregated.

3 · NIST AI RMF MAPPING

How the Copilot maps to the NIST AI Risk Management Framework.

The NIST AI Risk Management Framework (AI RMF 1.0, January 2023) is a voluntary framework organised around four functions — **Govern**, **Map**, **Measure** and **Manage** — and seven characteristics of trustworthy AI. NIST's Generative AI Profile (NIST-AI-600-1, July 2024) extends it with actions specific to generative systems. The table below summarises how Veroxos operationalises each function for the Copilot.

FUNCTION	WHAT NIST ASKS	HOW VEROXOS IMPLEMENTS IT
Govern	Cultivate a risk-management culture: policies, roles, accountability structures and lifecycle ownership for AI systems.	AI Copilot governance sits inside the ISO 27001 ISMS: a named asset owner and data owner, a documented AI usage policy, risk-register entries for AI-specific risks, and change control through the same CAB process as any other production change.
Map	Establish context: what the system is for, who it affects, and where risks can arise across the lifecycle.	The Copilot's intended use is documented and bounded (query, explain, recommend, draft — within TEM/mobility/asset data). Foreseeable misuse (e.g. treating drafts as decisions, over-reliance on low-confidence answers) is identified, with mitigations designed in: confidence signalling, evidence citation, and mandatory approval gates.
Measure	Assess and track AI risks with appropriate methods and metrics, including performance, reliability and unintended behaviour.	Answer accuracy is measured against reconciled ground truth (the deterministic reconciliation engine that underlies the platform); anomaly precision/recall is tracked per detection category; user rejection rates of drafted actions are monitored as a quality signal; regressions are reviewed before model or prompt changes ship.
Manage	Prioritise and act on risks: allocate resources, respond to incidents, and communicate with affected parties.	AI-related incidents follow the ISMS incident-response process with an AI-specific runbook; identified risks are prioritised in the risk register with owners and review dates; customers are notified of material AI-behaviour changes through release notes and, where contractually agreed, direct notice.

Against NIST's trustworthy-AI characteristics, the design choices in this paper map directly: **accountable & transparent** (audit trail, evidence citation), **explainable** (answers cite records), **privacy-enhanced** (no-training default, least privilege), **secure & resilient** (ISMS controls, segregation), and **valid & reliable** (measurement against reconciled ground truth).

4 · EU AI ACT READINESS

Where the Copilot sits under the EU AI Act, and what we are doing about it.

The EU AI Act (Regulation (EU) 2024/1689) entered into force on 1 August 2024 and applies in stages: prohibitions and AI-literacy obligations from 2 February 2025, general-purpose AI model obligations from 2 August 2025, and the main obligations for high-risk systems phasing in from 2 August 2026, with certain embedded-product categories following in 2027. The Act classifies systems by risk: unacceptable (prohibited), high, limited (transparency obligations), and minimal.

Our classification assessment

The Veroxos AI Copilot is an enterprise decision-support tool operating on commercial telecom, mobility and asset data. It is not used for any Annex III high-risk purpose (it makes no decisions about individuals' access to employment, credit, essential services, education, law enforcement or similar), and it performs no prohibited practice. Our assessment, kept under review, is that the Copilot falls under **limited-risk transparency obligations**: users must be aware they are interacting with an AI system, and AI-generated content must be identifiable as such.

What that means in practice

- → **Transparency (Article 50):** the Copilot is visibly labelled as AI in the interface; drafted outputs are marked as AI-drafted until a human approves them; exported AI-generated reports carry a generation notice.
- → **Provider/deployer roles:** StableLogic acts as provider of the Copilot feature; customers are deployers. We supply the documentation deployers need — intended purpose, capabilities, limitations, and the human-oversight design in this paper — to support their own AI Act obligations, including AI-literacy duties for staff who use the system.
- → **Upstream model obligations:** where the Copilot uses third-party general-purpose models, we track those providers' Article 53 documentation (technical documentation, copyright policy, training-content summary) and reflect material changes in our own assessments.
- → **Reassessment trigger:** if the Copilot's capabilities ever expand toward an Annex III purpose, classification is reassessed *before* release through the ISMS change-control gate — not after.

Status label: transparency measures — implemented. Deployer documentation pack — implemented, updated per release. Upstream GPAI documentation tracking — in progress as provider documentation matures across the industry. Formal reassessment cadence — annual, or on capability change.

5 · AUDIT-LOG ARCHITECTURE

“We have an audit log” is not an answer. Here is the schema.

Every Copilot interaction generates an immutable event record. Events are append-only, integrity-protected, retained separately from operational data, and available to customer administrators for export. The event families and core fields:

EVENT FAMILY	RECORDED WHEN…	CORE FIELDS
ai.query	A user asks the Copilot a question and receives an answer.	`event_id` `timestamp` `actor_id` `role` `query_text` `data_scopes` `evidence_refs` `confidence` `model_version`
ai.anomaly	The anomaly engine flags a deviation (rate mismatch, usage spike, zero-usage line, feature drift).	`anomaly_type` `detection_rule` `affected_records` `estimated_value` `evidence_refs`
ai.recommendation	The Copilot proposes an action (dispute, optimisation, disconnection, report).	`recommendation_id` `basis` `draft_content_ref` `estimated_impact` `requires_approval`
ai.approval	A human approves, amends or rejects a drafted action.	`recommendation_id` `approver_id` `decision` `amendments` `decision_timestamp`
ai.execution	An approved action is carried out by the platform.	`recommendation_id` `executed_by` `result_status` `downstream_refs`
ai.config	AI-relevant configuration changes: training opt-in status, model version rollout, approval-policy change.	`change_type` `old_value` `new_value` `changed_by` `change_ticket_ref`

Retention, access and integrity

- → **Retention:** AI audit events are retained for the life of the customer engagement plus the contractually agreed post-termination period (default 12 months), unless the customer specifies longer to meet its own regulatory obligations.
- → **Access:** customer administrators can view and export their organisation’s AI audit events; StableLogic access is role-restricted and itself logged.
- → **Integrity:** events are append-only with sequence integrity checks; deletion is not exposed as a product function. Corrections are made by compensating events, preserving the original record.

6 · DATA HANDLING & THE NO-TRAINING OPT-IN

Your invoice data is your invoice data. Training is opt-in, scoped, and reversible.

Default posture

- → Customer data (invoices, contracts, inventory, usage, tickets, disposition records) is **not used to train or fine-tune foundation models**. This is the default for every engagement and requires no configuration.
- → Prompts and retrieved context sent to any underlying model are used to generate the response, not retained for provider-side training; provider agreements are selected and reviewed on this basis.
- → Data residency and processing follow the engagement's contractual terms; GDPR roles (controller / processor) are set out in the data processing agreement, with cross-border transfer mechanisms documented where applicable.
- → Operational telemetry that contains no customer content (latency, error rates, feature usage counts) is used to run and improve the service — this is service telemetry, not model training.

The opt-in workflow, end to end

1. **Contract-level election.** Opt-in is agreed in writing at contract level by an authorised customer signatory — never via an in-product toggle clicked by an end user, and never bundled into general terms.
2. **Scope definition.** The election names exactly which data classes are in scope (e.g. anonymised anomaly patterns), the purpose (improving detection quality), and what is excluded (identifiable rates, contract terms, personal data).
3. **De-identification gate.** In-scope data passes an anonymisation/aggregation step before any training use; records that cannot be reliably de-identified are dropped, not "best-efforted."
4. **Logged activation.** The opt-in status change is recorded as an **ai.config** audit event with the contract reference, visible to customer administrators.
5. **Withdrawal.** Opt-in can be withdrawn with contractual notice; withdrawal stops future training use of the customer's data. The withdrawal is likewise logged as an **ai.config** event.

Straight answer to the common question: "Does Veroxos train AI on our data?" — No, unless you have signed an explicit, scoped opt-in, and you can see in your own audit log whether that has ever happened.

7 · HUMAN-IN-THE-LOOP APPROVAL DESIGN

Autonomy is a dial, and consequential actions sit behind a human gate.

The Copilot's outputs are tiered by consequence. The tiers, and who can approve what, are policy objects the customer controls — changes to them are themselves logged (**ai.config**).

TIER	EXAMPLES	GATE
Tier 0 — Answer	Query responses, explanations, report views, benchmark comparisons.	No approval required. Logged (**ai.query**) with evidence references and confidence.
Tier 1 — Draft	Drafted carrier dispute, drafted optimisation recommendation, drafted customer-facing report.	Requires named-user approval before anything leaves the platform or changes state. Draft is marked AI-generated until approved.
Tier 2 — State change	Plan change requests, line disconnection requests, inventory status changes with financial effect.	Requires approval by a user holding the specific role permission for that action class; optional second-approver rule for actions above a customer-set value threshold.
Tier 3 — Excluded	Payments, contract execution, access-control changes, anything affecting individuals' rights.	Not available to the Copilot at all — excluded by design, not by policy setting.

Design details that make oversight real, not ceremonial

- → **Evidence at the point of decision.** The approval screen shows the draft *and* its basis — invoice lines, contract clause, computed impact — so approval is informed, not a rubber stamp.
- → **Amend, don't just accept/reject.** Approvers can edit a draft before approval; the amendment is recorded, and systematic amendments feed the quality metrics in Section 3 (Measure).
- → **No approval fatigue by design.** Tier 0 answers don't generate approval requests, keeping the queue meaningful. Rejection rates are monitored: a rising rejection rate is treated as a model-quality signal, not a user problem.
- → **Break-glass is absent on purpose.** There is no "auto-approve all" switch. Autonomy increases only by explicit policy change per action class, logged and reversible.

8 · FOR YOUR DUE-DILIGENCE FILE

Quick answers for security, legal and procurement reviewers.

QUESTION	ANSWER
Is customer data used to train AI models?	No by default. Only under a written, scoped, withdrawable contract-level opt-in (Section 6), with activation visible in the customer's own audit log.
Can the AI act without a human?	Read-only answers, yes (logged). Anything that leaves the platform or changes state requires named-user approval; payments and rights-affecting actions are excluded entirely (Section 7).
What is logged, and can we export it?	Six event families covering query, anomaly, recommendation, approval, execution and configuration, with the fields listed in Section 5. Customer administrators can export their events.
How long are AI audit logs retained?	Life of engagement plus 12 months by default; longer by agreement.
Where does the Copilot sit under the EU AI Act?	Assessed as limited-risk (transparency obligations), implemented per Article 50; reassessed annually and on capability change (Section 4).
Is there a NIST AI RMF certification?	No such certification exists — the RMF is voluntary. Section 3 shows how each RMF function is operationalised so you can evaluate substance rather than badges.
What underlying security framework applies?	ISO 27001-certified ISMS (StableLogic); SOC 2 Type 2 audit in progress, expected FY28; HIPAA-ready and GDPR-compliant design; NIST SP 800-88 for media sanitisation.
Can we see this in practice before contracting?	Yes — the demo includes the approval workflow and a live view of the audit events your own questions generate.

Document control

This paper is reviewed at least annually and on any material change to the Copilot's capabilities, model suppliers, or the regulatory position. Feedback and due-diligence questionnaires are welcome via veroxos.com/contact (choose "Security / vulnerability disclosure" for security matters).

◆ VEROXOS

See the governance, live.

Bring your hardest due-diligence question to a 30-minute session. Ask the Copilot something on your own numbers — then watch the audit events it generates and the approval gate it respects.

REQUEST A DEMO →