Tucker Payment Processing · Standards self-assessment

ISO/IEC 42001:2023 AI Management System

The management system behind Tucker, clause by clause and control by control · Self-assessment · November 24, 2026

13 met **3 partly met**

About this document

This is a self-assessment written by Michael Cunningham, who designed and built the site. It is not a certification: no accredited body has audited it. ISO/IEC 42001 is the international standard for an AI management system (AIMS): the policies, roles, risk assessments, controls and reviews an organization uses to run AI responsibly. Clauses 4 to 10 are the management system itself; Annex A lists the reference controls. This document shows how a one-person fictional company would meet each, with evidence, and says plainly where a real organization would need more.

Tucker Payment Processing's role under the standard: it is the provider of two AI systems (Tucker, the support agent, and the AI SOC Squirrel, the security analyst), built on a general-purpose model supplied by Google (Gemini).

The management system (clauses 4 to 10)

Clause 4: Context of the organization — MET

The scope is defined: two AI systems on tuckerpaymentprocessing.com, serving anonymous public visitors with fictional data. Interested parties are named (visitors, the owner, the model supplier, hosting and security providers, and anyone reviewing the work), and so is what each expects: privacy, honesty about the AI, no harm from its actions.

Evidence: `README.md`, `policies/ai-assistant-policy.md`.

Clause 5: Leadership — MET

An AI policy exists, is published, and is owned by the top of the organization. Policy TPP-AI-01 sets what the AI may and may not do. The owner is accountable and is named on every page.

Evidence: `policies/ai-assistant-policy.md`, `site/policies.html`.

Clause 6: Planning — PARTLY MET

AI risks are identified and treated: each risk on the OWASP list is mapped to a control, and each AI's abilities are limited to what it needs. The impact on people is assessed (see A.5 below). Objectives are measurable: spend caps, rate limits and zero full card numbers in any output. A certifiable system would keep a formal risk register and a Statement of Applicability, which this set of write-ups stands in for.

Evidence: `docs/standards/owasp-llm-top-10.md`, `app/config.py`.

Clause 7: Support — MET

Resources are sized and capped (one server, a fixed AI budget). Competence is the owner's 25 years in information security and payments, plus AI-specific study. Communication with visitors is built into the site (Behind the Scenes, the About page), and documented information is under version control in a public repository.

Evidence: `README.md`, `deploy/README.md`.

Clause 8: Operation — MET

Operational controls are code, so they run the same way every time: identity checks in Python, read-only tools, output masking, limits and caps. Changes go through GitHub, with CodeQL, secret scanning and Dependabot checks, and are deployed with one scripted command that waits for any attack run in progress to finish.

Evidence: `app/chat.py`, `app/tools.py`, `deploy/update.sh`.

Clause 9: Performance evaluation

PARTLY MET

Monitoring is continuous: a nightly security and activity report, real-time alerts, more than 70 automated tests and live attack runs. Independent scanners checked the site on November 24, 2026. A certifiable system would add scheduled internal audits by someone other than the builder, and formal management reviews.

Evidence: `app/report.py`, `app/alerts.py`, `tests/`, `docs/scans/`.

Clause 10: Improvement

MET

Problems are fixed and prevented from coming back. Every gap the independent scanners found was fixed the same day. When a deploy interrupted an attack run, the cause was fixed and a test was added so it cannot recur. The commit history is the corrective action log.

Evidence: `docs/scans/2026-11-24/results.json`, `tests/test_cipher_soc.py`.

Annex A controls

A.2 Policies related to AI

MET

A documented AI policy (TPP-AI-01) that fits the business, alongside the related card security, refunds and data retention policies. The AI answers from these same documents, and a test keeps the published copies identical.

Evidence: `policies/`, `tests/test_policy_pages.py`.

A.3 Internal organization

PARTLY MET

Roles are clear for a one-person company: the owner holds every role and receives every report and alert, and vulnerabilities can be reported privately. A real company would split these roles (AI owner, security, privacy, an independent reviewer) and document who decides what.

Evidence: `SECURITY.md`, `app/alerts.py`.

A.4 Resources for AI systems

MET

Every resource the AI depends on is documented: the model and its version setting, the data it can read, the tools, the server, the database roles and the hosting and security services.

Evidence: `app/config.py`, `db/roles.sql`, `deploy/README.md`.

A.5 Assessing impacts of AI systems

MET

The impacts on people were assessed and designed against: privacy (data only for a verified customer, minimized), financial harm (no tool can move money) and deception (Tucker discloses that he is an AI). Neither AI makes a decision about a person, such as credit, eligibility or fraud. A real company would document this in the format of ISO/IEC 42005:2025, the companion standard for AI system impact assessments published in May 2025.

Evidence: `app/tools.py`, `app/session.py`, `app/chat.py`.

A.6 AI system life cycle

MET

Requirements were written as policy first, then built, tested and deployed with records at each step: design notes in the code, automated tests (the AI is replaced by a scripted stand-in so tests are repeatable), version-controlled deployment, and operation with monitoring.

Evidence: `app/llm.py`, `tests/`, `deploy/update.sh`, `deploy/README.md`.

A.7 Data for AI systems

MET

The data the AI touches is known, minimal and fictional: demo customers, cards (last four digits only), orders and four policy documents. Its origin is documented, and it is reset to a known state every night. Chat text is not stored with IP addresses, and web logs are kept for at most two days.

Evidence: `db/schema.sql`, `db/seed.sql`, `db/nightly_reset.sql`, `deploy/logrotate-apache2`, `policies/data-retention.md`.

A.8 Information for interested parties

MET

Visitors are told what the AI is, what it can and cannot do, and what happens to their data: the AI disclosure in each chat, the AI Policies pages, the About page, the privacy notes and the live Behind the Scenes panel. Security researchers have a published reporting channel.

Evidence: `site/about.html`, `site/policies.html`, `site/.well-known/security.txt`.

A.9 Use of AI systems

MET

Each AI is limited to its intended use. Tucker stays on Tucker Payment Processing topics and politely declines everything else. Each chat page gets only the tools it needs. The AI SOC Squirrel only ever sees one attack run's events.

Evidence: `app/chat.py`, `tests/test_page_chats.py`, `app/cipher/soc.py`.

A.10 Third-party and customer relationships

MET

Supplier responsibilities are understood and bounded. Google supplies the model on its paid tier (no training on the data) and receives no IP addresses or full card data. Cloudflare and Google Cloud host and protect the site. Every supplier is listed on the About page.

Evidence: `site/about.html`, `app/llm.py`.

What certification would add

A certifiable AI management system would add a formal AI risk register and Statement of Applicability, documented internal audits and management reviews, competence records and training plans for each role, supplier assessments, and an independent certification audit. The controls themselves are already working and tested here; what's missing is the paperwork and independence a larger organization needs.

Source: https://github.com/mcunningham-infosec/tucker-payment-processing/blob/main/docs/standards/iso-42001.md · Live site: https://tuckerpaymentprocessing.com
Written by Michael Cunningham with AI assistance (Claude). A self-assessment, not an audit, certification or legal opinion. Tucker Payment Processing is a fictional company and processes no real payments.