FOR GENERAL COUNSEL AND CHIEF FINANCIAL OFFICERS

Why the standards basis behind Darwin is different

When a health system adopts AI governed under Darwin, the governance and evidence infrastructure for consequential AI and agent decisions in healthcare, its governance rests on standards developed through accredited standards-development processes, not on vendor frameworks. The distinction is procedural, and procedure is exactly what makes a basis defensible: accredited processes require openness, balance of interests, due process, public review, consensus, and appeal, which produce a rigor and traceability that no white paper, alliance manifesto, or proprietary framework can claim. When governance is later examined, by a board, an underwriter, a rating agency, or opposing counsel, the question is always what the requirement rested on. This one rests on process designed to survive that examination.

The standards. IEEE/UL 2933:2024, developed through the IEEE Standards Association process, addresses trust, identity, privacy, protection, safety, and security for clinical systems; Darwin runs in production as its reference implementation. ANSI/HSI 2800:2025, the Hospital AI Operations Governance Standard, is an American National Standard developed under ANSI's essential requirements of due process and consensus. Medigram's CEO co-chairs the IEEE/UL 2933 Trust Subgroup and led the TTIC contribution to the authorship of ANSI/HSI 2800:2025: named roles in the public record, checkable at the source.

The published record. The written body of work behind this practice is governed the way scholarship is governed, not the way marketing is: the accompanying book series at Taylor & Francis is edited under Routledge's published Series Editor ethics statement, the editorial ethics framework of academic publishing, with disclosed conflicts and accountable editorship. That record is load-bearing, not decorative: the books were used in the accreditation review cycles for the standards, meaning the published literature and the accredited standards each stood behind the other's review.

From paper to evidence. Standards become defensible only when implemented and checked: Darwin's production behavior is independently verified against its documentation, holding the highest recorded behavioral verification score, as documented by the Open Agent and AI Security Community's independent publication, and produces a Governed Decision Record, sealed at the moment of every agent action, evidence that exists before anyone asks for it.

What this means for treasury and legal. This chain, accredited-process standards, named authorship, a production reference implementation, independent behavioral verification, and a governed published record, is documented governance evidence in the language rating agencies, bond investors, and D&O underwriters already trust. It supports enterprise risk narratives, rating-agency reviews, bond disclosures, and insurance conversations by answering the governance question with citable process rather than vendor assurance. The result, in the terms that matter to treasury and legal: know what was authorized, verify what happened, preserve the evidence, know who is accountable.

> *The Trustworthy Technology & Innovation Consortium (TTIC) is a non-commercial standards body; it does not endorse, attest, or certify Medigram. Though Medigram leads by example, operating the verification and assessment practice ahead of TTIC formalizing certification.*

REFERENCES

Medigram Evidence · medigram.com/evidence
IEEE/UL 2933:2024 · standards.ieee.org/ieee/2933/7592/
ANSI/HSI 2800:2025 · webstore.ansi.org/standards/ansi/ansihsi28002025
Routledge Series Editor ethics statement · routledge.com/rsc/downloads/SeriesEditor_Ethicsstatement.pdf
Independent verification publication · open-agent-ai-security.github.io/blog/ciso-defensive-captain-medigram/