

Dear IEEE Security & Privacy Magazine editors,

Please consider “Brokered Trust for AI Agents: Lessons from the eIDAS QTSP-Broker Pattern for an Agent Trust Framework” as a Regular Feature Article for IEEE Security & Privacy Magazine. The article is a practitioner-facing feature proposing the QTSP-broker pattern as architectural prior art for AI-agent trust frameworks. It targets the magazine’s practitioner audience: PKI engineers, AI architects and compliance leads who deploy or evaluate agent-trust systems.

The contribution is threefold. First, the single-root-of-trust assumption common in agent-trust proposals fails in real ecosystems. Second, QTSP-brokers under eIDAS already solved federation for humans by hiding many trust providers behind one operational interface. Third, the pattern maps onto open agent-trust design questions, with four concrete recommendations for the EATF reference framework.

Related but non-overlapping work is under review or available as preprint. BSEADE Paper A on agent-ecosystem accountability is under review at EJRR with a different framing and audience. EATF/AEP attestation work is documented in the IACR ePrint 2026/109612 preprint (non-peer-reviewed) and in the ACM CSUR survey submission `sokolov2026csur`; both address a different scope and layer. The present article is on the identity-broker layer and does not overlap with the attestation-evidence layer covered by the IACR ePrint/CSUR work or by OVERT 1.0.

Generative-AI assistance. This article was prepared with mixed-source generative-AI assistance. The body prose was authored directly by the author across April–May 2026 (source of record: `outlines/brokered-trust-for-ai-agents-v0.2.md`). The IEEE Security & Privacy Magazine submission packet — the 50-word abstract distillation, the formatted bibliography, and the cover letter — was prepared by an autonomous AI worker (Hermes CLI orchestrating OpenAI Codex / GPT-5.5) from a structured brief authored by Anthropic Claude Code (Opus 4.7) in dialogue with the author. All substantive claims, results, design decisions, and references are the author’s own and verified by the author, who takes full responsibility. Reported per COPE and ICMJE recommendations. IEEE Author Portal has a separate AI-use field; the same text should go there.

Anton Sokolov Researcher, Tyche Institute, Tallinn, Estonia `anton.sokolov@tyche.institute` ORCID: 0000-0003-2452-7096