JOHN M. WILLIS

PRINCIPAL AI GOVERNANCE ARCHITECT & TECHNICAL ASSURANCE ADVISOR

john@johnmwillis.me | +1 (240) 720-7515 | johnmwillis.me | linkedin.com/in/johnmwillis

github.com/jwillisSFT/agcp-spec | github.com/jwillisSFT/pbsai-reference-architecture | ORCID 0009-0007-5889-9628

PROFESSIONAL PROFILE

Independent technical assurance architect focused on what evidence actually establishes about AI systems: specification conformance, governance admissibility, engineering integration, security-control effectiveness and operational readiness. Creator of RGE and AGCP, with extensive experience assessing complex software, security architectures, authorization packages and failed production implementations.

DISTINCTIVE EXPERTISE

- Requirements-based, reproducible capability assessments and structured conformance evaluations.
- Semantic equivalence, negative testing, evidence provenance, test adequacy and engineering decision records.
- AI agent execution governance: deterministic authorization, admissibility, current-state and invariant checks.
- Independent architecture review, RMF/ATO assessment, software assurance, root-cause analysis and release verification.

CURRENT ARCHITECTURE & RESEARCH

SUSTAINABLE FUTURE TECH, INC. | Founder & Chief Innovation Officer; Principal Architect | Aug 2024–Present

- Developed AGCP conformance and evidence-based assessment methodology, separating open specifications from private evaluation and implementation.
- Created controlled AI-assisted SDLC methods distinguishing generated outputs from accepted changes: canonical-source integration, regression testing, diagnosis and closure evidence.
- Defined assessment concerns for unsupported completion claims, partial fixes, schema-only success and gaps between documentation and observed behavior.
- Sponsored the 2026 CMU Heinz capstone, providing requirements and architectural review for trusted-runner governance integration.

TURNAROUND SECURITY | Founder; C2C Security Architecture Consultant | Nov 2019–Mar 2025

- Delivered independent enterprise and federal security architecture work through a C2C contracting entity; pursued security-control modeling and AI-assisted engineering research.

SELECTED FEDERAL ENGAGEMENTS

- **Library of Congress / Congressional Research Service — via Zermount (Jul–Sep 2020):** Investigated a critical PIV authentication failure, identified certificate-authority server issues and delivered root-cause and mitigation recommendations.
- **USPTO — via Zermount (Oct 2022–Feb 2025):** Security/cloud architecture and HVA reviews, NIST/FedRAMP-aligned controls, AWS designs and capacity/performance alternatives.
- **DHS Science & Technology — via ICS (Aug 2013–Jan 2014):** Chief Security Specialist; directed security assessments and reviewed ATO packages before CIO/CISO decisions, including Fortify and Nessus technical findings.
- **TSA — via Zermount (2021–Feb 2025):** GRC task order and security control assessment leadership; Zero Trust identity advice and control modernization.

EARLIER TECHNICAL & LEADERSHIP EXPERIENCE

- **CMS — via CSC/GDIT (Oct 2014–Oct 2015):** Principal ISSE/Chief Security Architect; analyzed SALARS requirements, data flows and missing deliverables to guide recovery and production integration.
- **U.S. Mint — via SupremeSoft/Lockheed Martin (Dec 2011–Dec 2012):** Directed Nessus/AppScan vulnerability remediation; reversed failed Cisco NAC rollout; supported security maturity assessment and changes.
- **World Bank Group — via Experis (Dec 2010–Dec 2011):** Developed and updated secure coding, system configuration and privacy-framework reference materials.
- **CA Technologies / Layer 7 — employee (Feb–Sep 2014):** Consolidated fragmented engineering knowledge and documented best practices for REST/API and Mobile Access Gateway solutions.
- **Release and configuration assurance:** USAA WebSphere baseline processing reduced from approximately eight hours to 30 minutes; MUOS build/installation traceability; earlier enterprise CM, release qualification and production change control.

EDUCATION & PROFESSIONAL DEVELOPMENT

- **University of Sussex:** MSc with Merit, Sustainable Development (degree conferred 2023).
- **MIT Professional Education:** Applied Data Science Program (2024); built and compared CNN classifiers for malaria blood-smear imagery. Selected model achieved 98.27% accuracy on a 2,600-image held-out test set (academic capstone; not clinically validated).
- **Oxford:** AI Governance: Technology, Ethics and Institutions (2024); Senior Leadership Executive Program (2024).
- **Carnegie Mellon University:** CISO Executive Education and Certificate Program (2018).
- **ISC2 credentials earned:** CISSP (2008), CISSP-ISSAP (2016), CSSLP (2015), CCSP (2017).

PUBLICATIONS, EVIDENCE & TECHNICAL PORTFOLIO

- **Author:** *Runtime Governance Engineering*, Vol. 1: Foundations and Architecture (2026), ISBN 979-8-9971149-0-9; Vol. 2: Secure-by-Design AI—Deterministic Runtime Governance (2026), ISBN 979-8-9971149-2-3.
- **Author:** *Hybrid Quantum AI Systems Engineering* (2026), ISBN 979-8-9971149-1-6.
- **Research:** Governance compilation, formal execution semantics and safety invariants, AI-assisted SDLC assurance, and sequential AI inference.

PUBLIC TECHNICAL MATERIALS

AGCP open specification: github.com/jwillisSFT/agcp-spec
PBSAI reference architecture: github.com/jwillisSFT/pbsai-reference-architecture
Research & publications: orcid.org/0009-0007-5889-9628