

# Owen H. Eskew, MBA, MSAI

Enterprise architect and AI builder. Twenty-five years across USAF operations, DoD enterprise architecture, and a working AI product studio, with production systems, paying customers, published research, and a published book as verification.

New Braunfels, TX · owen.eskew@onsq.net · +1 (830) 302-3330, answered by his own AI receptionist · onsq.net · LinkedIn · Active Top Secret clearance

## AT A GLANCE The record in numbers

|        |                                                                                                                   |
|--------|-------------------------------------------------------------------------------------------------------------------|
| 10+    | AI systems designed, built, and operated; four in production with real users                                      |
| 25 yrs | in IT and data systems, from USAF operations through DoD enterprise architecture to founder                       |
| 200+   | DoDAF-compliant architecture artifacts delivered in DoD environments                                              |
| 755K   | Airmen served by the USAF learning systems he architected                                                         |
| 6      | LLM platforms in shipped work: Anthropic, OpenAI, Google, xAI, Microsoft Copilot, and custom endpoints            |
| 10+    | certifications, including CompTIA Security+ ce, Azure AZ-900 and AI-900, ITIL, GitLab, Google PM, and DoDAF/TOGAF |
| 3      | books published, six authored AI curricula, and an ACM SIGCSE 2026 submission                                     |
| TS     | Active Top Secret clearance; 25 years in Department of Defense environments                                       |

*"Advice is easy to sell and hard to verify. So here is the verification: AI products designed, built, and operated, with paying customers, live clients, students, and readers using them today."*

## LIVE PROOF Verify before you interview

Every major claim in this dossier has a public artifact. Call the phone number above and an AI system he built answers. The rest:

| ARTIFACT                              | WHAT IT PROVES                                                                           |
|---------------------------------------|------------------------------------------------------------------------------------------|
| <a href="#">ailura.net</a>            | Ailura AI Receptionist, a production voice/SMS/chat agent platform with paying customers |
| <a href="#">getailura.com</a>         | Ailura Home Services, a multi-tenant field-service SaaS running a live HVAC company      |
| <a href="#">cheqer.vercel.app</a>     | Cheqer, an AI Bible word-study app: hybrid RAG with published retrieval-quality metrics  |
| <a href="#">onsq.net/portfolio</a>    | Full product portfolio with the status of each system                                    |
| <a href="#">onsq.net/capabilities</a> | Capability statements: commercial, federal, education, speaking                          |
| <a href="#">Amazon author page</a>    | Published books, including <i>Image in the Machine</i> (2026)                            |
| <a href="#">store.onsq.net</a>        | ONSQ Enterprise Press, a working publishing and e-commerce operation                     |

EXPERIENCE

Professional experience

ServiceNow Business Process Analyst, AFLCMC/HNCD, U.S. Air Force

Feb – Sep 2026

Centurion (Leidos program). Cyber defense capability delivery on ServiceNow.

- ServiceNow ITAM, HAM, and SAM for an active acquisition directorate: asset tracking, hardware and software accountability, lifecycle reporting, and CMDB.
- Authored the directorate's ServiceNow user guide (Zurich release) covering ITSM, HAM, SAM, CMDB, SPM, and Change Management; built UAT test plans and requirements backlogs.
- Translated operational cyber needs into scoped, testable workflow and integration designs under accreditation and contract constraints.

Platform Engineer, Army Software Factory

Mar – Oct 2025

Strategic Intelligence Inc. (Akima), Austin, TX.

- Multi-cloud architectures integrating AWS GovCloud, Azure AI Foundry, and Palantir Foundry for enterprise data and AI workflows, plus AI automation frameworks and custom LLM work.
- CI/CD with GitLab, Terraform, and Kubernetes; coached Army soldier-engineers on cloud platforms, observability, and DevSecOps.

## Site Lead and Enterprise Learning Architect, HQ AETC A3G

Oct 2023 – Feb 2025

*Strategic Technology Institute, Randolph AFB, TX. Enterprise Architect for Air Force learning services.*

- **Integrated an AI chatbot into the USAF myLearning platform**, the Air Force's enterprise learning system, including conversational experience design, an output-analysis process, and the command's AI adoption strategy.
- **Led delivery of Power Automate service-management workflows for an Air Force major command directorate**: request intake, routing, approvals, and reporting built on the Microsoft Power Platform. Ran Microsoft Copilot in production from initial release, built in Copilot Studio, and authored the AI governance and acceptable-use framework covering agent approval, data-handling boundaries, and production review gates.
- Produced DoDAF-compliant architecture views, capability maps, data models, and system-integration diagrams for cross-domain training platforms; drove modernization integrating AI, microservices, API-based analytics, and cloud-native components into myLearning.
- Achieved the first-ever RMF Authority to Operate and USAF Program of Record for ITP(XR), an extended-reality maintenance-training platform.
- Site Lead responsible for 19 contractors; ServiceNow enterprise architecture (ITSM, ITPM, CMDB, governance); executive briefings to senior Air Force leadership.

## Enterprise Architect, USAF ADLS / TIMS / GTIMS

2014 – 2023

*Data Systems Analysts, Inc., JBSA Randolph, TX. Progressed from systems engineering and test leadership to Enterprise Architect and Infrastructure Team Lead.*

- Enterprise Architect for USAF training-management systems serving 755K+ Airmen across 81+ sites (GTIMS: 60K users, 80K workstations).
- Led GTIMS modernization to Azure GovCloud; stood up DevSecOps pipelines that cut deployment cycles 35%.
- Delivered 200+ DoDAF-compliant artifacts: system views, data flows, data models, integration patterns, and security models.
- Led ServiceNow implementation (ITSM, ITOM, CMDB); application-portfolio rationalization and shadow-IT consolidation; managed \$1.2M budgets; headed AETC's Service Validation and Testing service.
- Co-authored a 2019 AETC AI/ML pilot proposal: an ML scheduling engine fusing training data, weather feeds, and syllabi into risk-scored schedule recommendations with a Qlik leadership dashboard.

## Earlier: Leidos, Raytheon, U.S. Air Force

2002 – 2018

- **Leidos**: Senior System Engineer Manager for a data-center migration to AWS GovCloud; ServiceNow ITSM SME for CMDB, Change, and Incident Management. Earlier, depot engineering with predictive-failure metrics and trend analysis.
- **Raytheon**: cut a security accreditation process from over six months to four weeks; deployed an insider-threat monitoring system into USAF networks ahead of schedule.
- **USAF (SSgt)**: administered a network-warfare test range of 5,000 workstations; engineered 23 Unix, Solaris, and Windows data-analysis systems processing 34 TB per day at 99% uptime; led client systems for 3,000+ machines. AF Achievement Medal with Valor.

## Founder and Principal, ONSQ Enterprises / Ailura

2025 – present

*AI product studio, publisher, and advisory. Also CTO of Doctor Marriage / Mapping Manhood.*

- Designed, built, deployed, and operates every system in the portfolio below, including sales, support, compliance, and app-store distribution.

## Education and certifications

- **M.S. Artificial Intelligence** (Data Analytics), The University of Texas at San Antonio, 2026. **MBA**, Technology Management. **B.S.**, Information Technology. **A.S.**, Electronic Systems Technology (CCAF).
- **Certifications**: CompTIA Security+ ce (current), CompTIA A+, Microsoft Azure Fundamentals (AZ-900), Azure AI Fundamentals (AI-900), ITIL Foundations, ITIL Intermediate (Service Design and Service Transition), GitLab Certified Git Associate (2025), Google Project Management, Red Hat OpenShift, Enterprise Architecture (FEAC), DoDAF/TOGAF, DoD Cyber Awareness (current).

Each system below was architected, coded, deployed, and is operated by Owen personally. The stacks listed are what actually runs in the code.

## **Ailura AI Receptionist** IN PRODUCTION · PAYING CUSTOMERS

Multi-tenant AI front office answering voice, SMS, chat, and email for businesses. A tool-using agent books appointments against live calendars, quotes prices, relays messages, and escalates to humans, across all channels from one engine. Emotion AI (53-signal prosody detection) automatically escalates distressed callers. Tiered voice caching controls LLM and TTS cost; per-tenant OAuth calendar integrations; TCPA and A2P 10DLC compliance built in as code; automated nightly activity reporting.

STACK OpenAI GPT-4o and Realtime API, agent tool-calling, Hume emotion AI, ElevenLabs, embeddings with RAG; Python, FastAPI, PostgreSQL, Redis, Twilio and Telnyx, Google and Microsoft OAuth; GCP Cloud Run, Fly.io.

## **Ailura Home Services** IN PRODUCTION · LIVE HVAC CLIENT

Multi-tenant field-service and e-commerce SaaS running a real HVAC company: missed call to AI text-back, conversational SMS booking agent, route-optimized dispatch, voice-dictated notes, AI-drafted Good/Better/Best quotes, text-to-pay, QuickBooks sync, and nightly KPI reports. Claude vision reads receipts and equipment nameplates into strict schemas with confidence gating. Explainable customer-opportunity scoring carries mandatory reasons and confidence on every row, with fairness guardrails coded in. A database exclusion constraint makes double-booking structurally impossible; the schema, not the prompt, is the final referee over the AI.

STACK Claude agents and vision, Whisper transcription, explainable scoring; Next.js 15, React 19, Expo/React Native, Supabase Postgres with row-level security, 23 edge functions; Stripe Connect and Tap to Pay, QuickBooks OAuth sync, Twilio A2P 10DLC; Vercel, EAS.

## **Chequer** LIVE · CHEQUER.VERCEL.APP

AI-powered Bible word-study platform and graduate research artifact. A multi-turn Claude agent with four retrieval tools answers questions over a corpus of 425,000+ word tokens from eight ancient-text sources, using hybrid retrieval: exact lemma SQL where scholarly tagging exists, pgvector semantic search where it does not, with a deterministic citation-verification layer that structurally prevents hallucinated references. Retrieval quality is measured, not assumed: a reproducible evaluation harness reports precision@10 of 0.772 against a keyword baseline, with published per-case failure analysis.

STACK Claude Opus 5 tool-use agent, OpenAI embeddings, pgvector hybrid RAG, IR evaluation (precision@k, recall), anti-hallucination guardrails; TypeScript, Expo, Supabase with 20 SQL RPC functions, Python ETL over eight corpora; Vercel.

## **Ailura Quant** DEPLOYED · LIVE FORWARD TEST

Autonomous options-research and paper-trading engine running around the clock in the cloud. Black-Scholes pricing, greeks, and implied-volatility solving implemented from first principles; intraday regime classification; multi-factor candidate scoring; vectorized backtesting; and a read-only brokerage integration that coaches real trades against codified risk rules. Engineered with production reliability discipline: dead-man switches, crash-safe state recovery, audit logging, and 146 automated tests.

STACK Quantitative modeling, backtesting and analytics, natural-language command interface; Python, FastAPI, pandas, numpy, SQLite with raw SQL, Alpaca and Schwab APIs; Docker, Fly.io.

## NeuroDriver RESEARCH · CLASSROOM-TESTED

Browser-based 3D "glass-box" driving simulator that teaches students how autonomous vehicles perceive and decide. Physically simulated LiDAR, thermal, camera, and audio sensors feed a hybrid AI, combining a TensorFlow.js deep Q-network, a safety behavior tree, and transparent utility scoring, whose every decision is inspectable live. Deployed with 97 middle-school students; the mixed-methods study is under review at ACM SIGCSE 2026.

**STACK** TensorFlow.js DQN, reinforcement learning, explainable-AI interface, sensor simulation; React, Three.js, react-three-fiber, Vite.

## Ailura WordPress product line SHIPPED COMMERCIAL PRODUCTS

Two commercial WordPress plugin families demonstrating end-to-end productization. Ailura Guru is a site chatbot with a five-provider LLM abstraction (Claude, GPT-4o, Gemini, Grok, and custom endpoints), encrypted key storage, knowledge-base retrieval, and CRM lead capture. Ailura Blog Writer (eight tagged releases, in production for a nationally distributed counseling-education platform) does deep-research generation with enforced JSON schemas, DALL·E 3 imagery, and automatic SEO integration. Licensing, free and pro editions, marketplace compliance, and release engineering included.

**STACK** Five-provider LLM routing, structured outputs, DALL·E 3; PHP, WordPress, MySQL, Freemius licensing, release automation.

## STRONGHOLD and Vigilant Spirit MOBILE · EXPLAINABLE AI

STRONGHOLD (Android, Google Play pipeline) is a recovery-aware AI strength coach that fuses wearable biometrics, including HRV baselines and workload ratios, with training data. Claude adjusts the plan through schema-constrained structured outputs that server-side guardrails re-validate: the model can adapt but never violate the safety envelope. Includes hand reverse-engineering of a rowing machine's Bluetooth protocol. Vigilant Spirit is an AI dream journal with genuine explainability, showing SHAP and LIME attributions to end users to demonstrate transparent consumer AI.

**STACK** Claude structured outputs, Zod-constrained decoding, SHAP, LIME, biometric feature engineering; TypeScript, Expo, SQLite with Drizzle, Health Connect, BLE; Express, PostgreSQL.

# Technology inventory

| DOMAIN              | TECHNOLOGIES                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LLM platforms       | Anthropic Claude (Messages API, vision, tool use, structured outputs); OpenAI (GPT-4o family, Realtime API, Whisper, DALL·E 3, embeddings); Google Gemini and Vertex AI; xAI Grok; custom and self-hosted endpoints; Azure AI Foundry; Microsoft Copilot and Copilot Studio                                                                                                                                          |
| Agentic AI          | Multi-turn tool-use agent loops; structured outputs (JSON Schema, Zod-compiled formats); cross-provider tool-schema translation; model fallback and cost-tiered routing; guardrails as code (database constraints, citation verification, bounded action spaces); per-call token and cost ledgers; prompt-injection defenses; rate limiting                                                                          |
| RAG & retrieval     | pgvector with ivfflat indexing; hybrid lexical and semantic retrieval; embedding pipelines; retrieval evaluation (precision@k, recall, lift over baseline); SQL-grounded context injection; vector-store content pipelines                                                                                                                                                                                           |
| ML & explainability | SHAP; LIME; deep Q-networks (TensorFlow.js); spiking neural networks (96.95% accuracy at roughly 49% event retention); scikit-learn, PyTorch, TensorFlow; statistical learning in R; SAS; feature engineering (rolling baselines, z-scores, workload ratios); quantitative modeling (Black-Scholes, implied volatility)                                                                                              |
| Power Platform      | Power Automate: automated service-management processes delivered for an Air Force major command directorate (intake, routing, approvals, reporting); Microsoft Copilot in production since initial release; Copilot Studio; Power BI; authored AI governance and acceptable-use frameworks for enterprise Copilot and agent adoption                                                                                 |
| Voice & telephony   | Twilio (Voice, SMS, TwiML, Media Streams, signature validation, A2P 10DLC); Telnyx; OpenAI Realtime over media streams; ElevenLabs including voice cloning; Hume AI TTS and emotion; Google Cloud TTS/STT. Hand-built pipelines, no low-code wrappers                                                                                                                                                                |
| Languages           | Python; TypeScript and JavaScript; SQL and PL/pgSQL; PHP; PowerShell; Bash; R; SAS; C#; HTML/CSS                                                                                                                                                                                                                                                                                                                     |
| Data engineering    | PostgreSQL; Supabase (RLS, auth, edge functions, cron, webhooks); SQLite; MySQL; Redis; Drizzle ORM; SQLAlchemy; schema design for multi-tenant systems of 37, 16, and 14 tables; exclusion constraints, GIN and ivfflat indexes, triggers, audit logs; multi-source ETL (XML, TSV, REST, and scraping across eight corpora); QuickBooks and wearable-data sync pipelines; idempotent pipeline design; pandas, numpy |
| Analytics & BI      | Power BI; Qlik Sense; SSRS; SAS Enterprise Guide; automated KPI pipelines with daily, weekly, and monthly rollups; custom production dashboards; model and retrieval evaluation metrics; mixed-methods study design                                                                                                                                                                                                  |

|                         |                                                                                                                                                                                                                                 |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cloud & DevOps          | Azure GovCloud; AWS GovCloud; GCP (Cloud Run, Cloud SQL); Vercel; Supabase; Fly.io; Docker; Kubernetes; Terraform; GitLab CI/CD; Expo EAS with App Store and Google Play submission; secrets management and envelope encryption |
| Enterprise & governance | ServiceNow (ITSM, ITAM/HAM/SAM, CMDB, ITOM, SPM, Change Management); DoDAF (200+ artifacts); TOGAF; RMF and ATO delivery; ITIL; Zero Trust; UAT planning; Sparx EA, Visio; compliance as code (TCPA, CAN-SPAM, GDPR, A2P 10DLC) |

ROLE FIT

## What this means for a data-systems role

| CORE DUTY                  | TRACK RECORD                                                                                                                                                                                            |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Requirements → data models | DoDAF data models and capability maps for Air Force enterprise systems; three production database schemas designed from scratch; ServiceNow requirements backlogs and workflow specifications at AFLCMC |
| SQL & databases            | Hand-written PL/pgSQL across 60+ migrations; stored-procedure API layers; row-level security design; constraint-driven data integrity; index and performance tuning                                     |
| ETL & pipelines            | Eight-source scholarly-corpus ETL; QuickBooks 15-minute sync; wearable-biometric aggregation; market-data ingestion; idempotency and auditability as standing design rules                              |
| Dashboards & reporting     | Power BI, Qlik Sense, and SSRS in Air Force settings; automated KPI reporting with rollups and live operations dashboards in his own production systems; SAS and R analytics training                   |
| Analysis & documentation   | Authored a full ServiceNow platform user guide for an AF directorate; 38 documentation guides across his own products; UAT plans; founded AETC's service validation and testing practice                |
| Stakeholder translation    | Executive briefings to Air Force senior leaders; site lead for 19 contractors; onboarded real pilot customers; a published author and speaker who explains AI to non-technical audiences for a living   |
| AI-augmented data systems  | His AI systems ship with evaluation harnesses, explainability columns, fairness guardrails, and cost ledgers: analyst discipline applied to generative AI                                               |

RESEARCH & WRITING

## Research, publications, and teaching

|                                                                                                                                                                                                                                                                                               |                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Research (UTSA, M.S. Artificial Intelligence)</b> <ul style="list-style-type: none"><li>NeuroDriver classroom study (97 students), under review at ACM SIGCSE 2026</li><li>Explainable AI: SHAP analyses on COMPAS and Communities-and-Crime; a dream-classification XAI project</li></ul> | <b>Publications and teaching</b> <ul style="list-style-type: none"><li><i>Image in the Machine</i> (2026), on AI, human identity, and meaning; book, companion journal, and a 12-session study curriculum</li><li>Two additional published books; founder of ONSQ Enterprise Press (Amazon, IngramSpark, WooCommerce)</li></ul> |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

- Spiking neural networks: 96.95% accuracy at roughly 49% event retention; a hyperdimensional associative memory architecture
- Cross-corpus retrieval evaluation with published precision and recall results (Chequer)
- SELFPRINT (in progress): metamorphic probing of LLM self-representation across six models, roughly 13,800 responses
- Six authored AI-literacy curricula spanning middle school through college and professional development
- Keynotes and workshops on AI adoption for schools, churches, and businesses; weekly essays on Substack