

# Jordan Truong — AI Engineer · Multi-Agent Systems & Infrastructure

**Location:** Glendale, AZ · On-site, remote, hybrid · Open to travel **Email:** [emailme@jordantruong.com](mailto:emailme@jordantruong.com) **Site:** [jordantruong.com](http://jordantruong.com) **GitHub:** [github.com/bludragon-ai](https://github.com/bludragon-ai) **LinkedIn:** [linkedin.com/in/jordan-truong-ai](https://linkedin.com/in/jordan-truong-ai)

---

## Professional Summary

AI Engineer building and operating production multi-agent systems. Maintainer of **all41n1411a** — an open-source MCP memory server for AI agents, published to PyPI and GitHub. Focuses on multi-model orchestration (Claude, Gemini, Perplexity), context architecture, failure-pattern diagnosis, and cost-optimized token economics. Solo-built two production operational tools (ACES Aid + Triage v2) deployed floor-wide across separate enterprise client campaigns at Valor BPO. Completed all 15 Anthropic Academy certifications (2026). 14 years across business ownership, enterprise IT, and cybersecurity training — now building autonomous systems that ship.

**Open to:** AI Platform Engineering, Infrastructure, and Automation roles.

---

## Technical Skills

**AI & Agents:** Multi-Agent Orchestration (Claude, Gemini, Perplexity) · Specification Design & Prompt Engineering · Context Architecture (RAG, window management, retrieval) · Cost & Token Economics · Model Context Protocol (MCP — servers, clients, transport, sampling) · Custom MCP Server Development · Structured Output / JSON Schema · Agentic Workflows · Tool Use · Subagent Orchestration · Agent Skill & Hook Authoring · Claude Code · LangChain · ChromaDB · FAISS · Vector Embeddings · Streaming · Vision/Multimodal

**Agent Development Stack:** Playwright · Chrome DevTools MCP (browser automation, performance audit) · Semgrep (static analysis, security scanning) · CodeRabbit (automated code review) · Vercel MCP (deployment automation) · Context7 (docs retrieval) · Claude Code plugin ecosystem (code review, testing, deployment, agent orchestration)

**Integrated Agent Workflows (MCP):** Gmail · Google Calendar · Google Drive · Notion · Obsidian · Stripe · Mercury · DocuSign · Clerk · Telegram — wired for autonomous cross-tool agent pipelines

**Development:** Python · JavaScript · TypeScript · React · Next.js · Node.js · Bash · REST APIs · HTML/CSS · Streamlit · Tailwind CSS

**Cloud & Infrastructure:** Amazon Bedrock · Google Cloud Vertex AI · Docker · Linux (Debian) · systemd · Vercel · Firebase · PostgreSQL · Redis · Caddy · Git · GitHub Actions · Stripe · CI/CD

**Security & Compliance:** HIPAA (PHI handling, audit-safe data flows, regulated-industry deployment) · PCI DSS · NIST SP 800-53 · NIST CSF · HMAC Auth · SSH Hardening · UFW · File Integrity Monitoring · Content Security Policy · Semgrep Static Analysis

**Operations:** ServiceNow · Zendesk · ITIL · SOP Development · Technical Documentation · KPI/SLA Analysis

---

## Certifications

**Anthropic Academy — All 17 Courses Completed (2026)**

Building with the Claude API · Introduction to Model Context Protocol · MCP: Advanced Topics · Claude Code in Action · Claude Code 101 · Introduction to Agent Skills · Introduction to Subagents · Introduction to Claude Cowork · Claude 101 · AI Capabilities and Limitations · AI Fluency: Framework & Foundations · Claude with Amazon Bedrock · Claude with Google Cloud Vertex AI · AI Fluency for Educators · AI Fluency for Students · Teaching AI Fluency · AI Fluency for Nonprofits

## Google AI Professional Certificate — Google Careers (Apr 2026)

Credential ID: 218d1556-dd7d-47f9-8411-28ed063e7c24

## Cybersecurity Consultant Training — BaxterClewis Consulting (2023–2024)

PCI DSS compliance · NIST SP 800-53 audits · NIST CSF implementation · risk analysis · security policy development

---

## Key Projects

### all41n141la — Open-Source MCP Memory Server for AI Agents [Maintainer · Published on PyPI + GitHub]

Portable long-term memory for AI agents. Markdown on disk; speaks Model Context Protocol. Solves the problem of LLM agents losing context between sessions — turns a local markdown filesystem into queryable memory accessible across any MCP-compatible harness.

- **Cross-harness compatibility:** wired into Claude Code, Claude Desktop, Codex, Gemini CLI, and Antigravity
- **Architecture:** Karpathy-style wiki layer (raw sources + LLM-maintained synthesis) over episodic daily logs and semantic clustering
- **MCP transport:** stdio + HTTP; OAuth 2.1 authentication path; structured tool schemas
- **Distribution:** `pip install all41n141la` (PyPI alpha) + source on GitHub under MIT
- **Tests:** 16/16 unit suites covering MCP transport, memory CRUD, search/recall
- **Tagline:** "Portable memory for AI agents. Markdown on your disk. Speaks MCP."

**Tech:** Python · Model Context Protocol · OAuth 2.1 · CLI · Markdown **Links:** [github.com/bludragon-ai/all41n141la](https://github.com/bludragon-ai/all41n141la) · [PyPI](https://pypi.org/project/all41n141la/)

### Valor BPO — Multi-Campaign Operational Tooling [Production · 2 Campaigns Shipped · In Daily Use]

Solo-built production tooling across enterprise client campaigns at a global Business Process Outsourcer. Repeatable pattern: shadow agents on live calls, identify the workflow gap, ship a single-file or bundled tool within client compliance constraints, deploy floor-wide via simple distribution.

#### Campaign A — T-Mobile (Wireless Telecom) · ACES Aid

- **162 troubleshooting flows** across 23 categories with dynamic field population
- **Fuzzy search** (Fuse.js) across 1,500+ indexed keywords and 90 aliases — bridges customer language and internal documentation
- **Template generator** auto-populates call documentation (case notes, escalation forms, incident reports, account records), eliminating copy-paste errors
- **Built-in tooling:** phone directory, callback scheduler, EOD report export, multi-timezone clocks, device and equipment tracking

- **Single HTML file · zero external dependencies · runs offline** — deliberate constraint matching restricted software-install policies and unreliable floor connectivity
- **Content Security Policy** headers, local-first data storage, zero external network calls
- **Adoption:** 40+ agents across all shifts, shared floor-wide via simple download
- **Outcome:** Presented to operations leadership; personal KPI bonus maxed via ACES Aid use; coached other agents who also hit max-bonus KPI tiers

**Tech:** Vanilla JavaScript · Fuse.js · Offline-first · Single-file architecture · Zero dependencies

### Campaign B — Healthcare Billing (RCM) · Triage v2

- **4 lines of business** with per-LOB brand inversion, payment rules, collections routing, and DOS-cutoff logic
- **72 verbatim KB-grounded scripts** across 11 categories · **27 structured RPM playbooks** with location markers and warn-out blocks · **41 canned-note codes** auto-suggested by call context
- **HIPAA-aware save flow:** PHI-bearing note routes to the rep's clipboard for paste into the system of record; only scrubbed metadata persists to disk — no PHI written to local storage
- **Coverage:** 102/102 caller-type × call-reason scenarios audited green · **21/21 unit-test suites passing** (templates, scripts, canned notes, playbooks, EOB calculator, CARC lookup)
- **Architecture:** modular ES module library ( `lib/*.js` ) bundled to a single offline HTML file via a custom build script; runs inside Obsidian Custom Frames on VDI
- **UX:** call-type pill menu (one-click loads scripts + playbook + canned note), guided-flow toggle, live note preview, KB-search side panel, ACES-v3 dual-accent theme

**Tech:** Vanilla JavaScript (ES modules) · Custom bundler · DOM API · Offline-first · 21-suite test runner

### Blu Print Engine — Autonomous Multi-Agent AI Platform [Active Development]

Self-hosted AI operations platform running on a Linux VPS. Multi-agent decomposition across specialized agents (Scout, Sentinel, Architect, Nimbus) handling market research, competitive intelligence, infrastructure monitoring, and strategic analysis.

- **Multi-model orchestration** with fallback-chain routing across Anthropic Claude, Google Gemini, and Perplexity Sonar Pro — model selected by task fit (Claude for reasoning, Gemini for throughput, Perplexity for real-time web)
- **Context architecture:** structured prompt specifications, token-aware chunking, model-specific window strategies
- **Trust boundaries & security:** HMAC-signed stateless auth, guardrails for autonomous agents, SSH hardening, UFW firewall, localhost-bound services
- **Infrastructure:** systemd services, cron scheduling, atomic JSON writes, watchdog monitoring, Syncthing cross-device sync
- **Spark Pipeline:** single Telegram command triggers AI market research, validation, code generation, auto-testing, Vercel deployment, and sales packaging — first product shipped in under 2 hours at \$0.03 per research cycle

**Tech:** Python · Claude API · Gemini API · Perplexity API · Linux · systemd · Telegram Bot · Syncthing

### SoloInvoice AI — AI-Powered Invoice Generator [Shipped]

First product shipped through the Spark Pipeline, idea to production in under 2 hours. Natural-language invoice creation, PDF export, and template management.

**Tech:** Next.js · Tailwind · shadcn/ui · AI Pipeline **Links:** [Live Site](#) · [GitHub](#)

## RAG Agent — Document Intelligence System [Open Source – Dockerized]

Context architecture over vector embeddings with multi-turn conversation memory, source-citation evaluation, and retrieval quality scoring. Retrieval scoring rejects low-confidence answers rather than presenting them.

**Tech:** LangChain · ChromaDB · Python · Streamlit · Docker **Link:** [github.com/bludragon-ai/rag-agent](https://github.com/bludragon-ai/rag-agent)

## SQL Agent — Natural Language Database Interface [Open Source – Dockerized]

Query any PostgreSQL database using plain English. SQL generation with injection prevention, read-only execution, output quality validation, and natural-language explanations. 20 unit tests.

**Tech:** LangChain · PostgreSQL · Python · Streamlit · Docker **Link:** [github.com/bludragon-ai/sql-agent](https://github.com/bludragon-ai/sql-agent)

## Email Agent — AI Email Triage & Automation [Open Source – Dockerized]

Autonomous email management — categorization, priority scoring, thread summarization, and smart-reply drafting with output evaluation loops. Pluggable provider architecture. 12 unit tests.

**Tech:** LangChain · Claude API · Python · Streamlit · Docker **Link:** [github.com/bludragon-ai/email-agent](https://github.com/bludragon-ai/email-agent)

---

## Professional Experience

### Founder · Blu Print Solutions, LLC

Phoenix, AZ · Apr 2026 – Present (*independent project work since Feb 2025*)

- Founded and operate an AI studio building autonomous multi-agent systems and AI-powered SaaS products
- Released **all41n141la** — an open-source MCP memory server for AI agents — to PyPI and GitHub
- Shipped Sololnvoice AI via an automated idea-to-production pipeline (live URL)
- Published three additional dockerized open-source agents (RAG, SQL, Email) on GitHub under `bludragon-ai`
- Stood up LLC, business banking, and self-hosted e-commerce (Stripe) ready for direct payment processing

### Central Technical Expert / Senior Agent — Multi-Campaign

Valor Global · Glendale, AZ · Aug 2024 – Present

- Solo-designed, built, and deployed operational tooling across two enterprise client campaigns (T-Mobile and Healthcare Billing); both in daily production use
- **Campaign A (T-Mobile):** built and deployed **ACES Aid** — 40+ agents, all shifts, enterprise wireless business support
- **Campaign B (Healthcare Billing):** built and deployed **Triage v2** — 4 lines of business, 72 KB-grounded scripts, 27 RPM playbooks, HIPAA-safe save flow, 21/21 test suites green
- Surfaced workflow gaps by shadowing agents on live calls; pitched solutions to ops management as levers for CRT, FCR, NPS, AHT, and QA scores
- Earned personal KPI recognition award for maxing bonus tiers directly attributable to ACES Aid usage
- Coached other agents on the floor into hitting max-bonus KPIs using the tools
- Serve as night-shift escalation authority for complex network, device, and account issues

- Established standardized procedures aligned with ITIL best practices, reducing tribal-knowledge dependency across teams

## Service Desk Support Analyst

**Message Envy (Corporate) · Scottsdale, AZ · Feb 2023 – Apr 2024**

- Managed provisioning and deprovisioning workflows across 1,000+ franchise locations via ServiceNow
- Led network and security audits with firmware upgrades across 1,000+ devices; maintained 99.8% uptime
- Authored standardized SOPs adopted company-wide, reducing support-ticket backlog and improving team onboarding
- Provided Level 1-2 support for POS systems, network infrastructure, and franchise IT operations

## IT Support Specialist II

**Alorica (Dell Technologies BPO) · Remote · May 2022 – Jan 2023**

- Delivered Level 2 technical support for Dell Technologies enterprise clients across desktop, Windows/macOS, mobile, and AV equipment via Zendesk; consistently met SLA targets
- Performed systematic root-cause analysis, improving resolution time across device, software, and connectivity categories
- Authored 70+ knowledge-base articles covering configuration procedures, testing protocols, and systematic troubleshooting methodology — increasing first-call resolution

## Independent Business Owner

**Self-Employed · Goodyear, AZ · Oct 2010 – Mar 2022**

- 12-year owner-operator (1099 sole proprietor) managing multi-site operations, staff, and client relationships
- Owned full P&L: bookkeeping, vendor negotiations, product and equipment procurement, inventory, and tax-ready records
- Ensured ongoing regulatory and state-board compliance across all locations
- Built and executed client-retention marketing — email and SMS campaigns for repeat customers, seasonal promotions, and holiday offers
- Managed end-to-end operations including scheduling, inventory, customer retention, and financial reporting

---

## About

I build production AI infrastructure — multi-agent platforms, LLM-powered automation, and tools that handle real workloads.

My background isn't typical. I spent 12 years running my own business before moving into enterprise IT, cybersecurity training, and now AI engineering. I've seen how real users break software, how networks fail, and how manual processes drain resources. I build for the messy reality, not the demo.

## How I Work

1. **Architect** — Design systems that run autonomously, not tools that need babysitting
  2. **Orchestrate** — Deploy AI agents across models and APIs to execute at scale
  3. **Ship** — Idea to deployed product in hours, not months. Live URLs, real users
-

## The 7 AI Skills That Matter in 2026

Skill	How I Prove It
<b>Specification Precision</b>	Structured prompt specifications across multi-model orchestration; prompt engineering across 162 operational workflows in ACES Aid
<b>Evaluation &amp; Quality Judgment</b>	Output quality validation across RAG, SQL, and Email agents; combined unit tests; retrieval quality scoring that rejects low-confidence answers
<b>Multi-Agent Decomposition</b>	4 specialized agents (Scout, Sentinel, Architect, Nimbus) decomposing business objectives into scheduled tasks; subagent orchestration
<b>Failure Pattern Recognition</b>	Structured fault diagnosis across API, model, and infrastructure layers; fallback-chain routing; watchdog monitoring
<b>Trust &amp; Security Design</b>	HMAC-signed auth, PCI DSS compliance, NIST frameworks, guardrails for autonomous operation
<b>Context Architecture</b>	RAG over vector embeddings, token-aware chunking, model-specific context window strategies, retrieval optimization
<b>Cost &amp; Token Economics</b>	\$0.03 per research cycle, deliberate token optimization, multi-model cost arbitrage across Claude / Gemini / Perplexity