

Spencer Presley

spencerpresley96@gmail.com | linkedin.com/in/spencerpresley96 | github.com/SpencerPresley

EXPERIENCE

CrunchAtlas — Software Engineer | Nov 2025 – Present

Local-first, air-gapped AI security platform for critical infrastructure. Sole author of the systems below.

- Built the **Postgres-native job queue** that replaced Celery for the platform's GPU, LLM, and sandbox jobs, so a burst of heavy AI work can't OOM an air-gapped box and any one job can be cancelled without killing its worker. It runs on atomic per-lane claiming (`SELECT FOR UPDATE SKIP LOCKED`), priority-FIFO ordering, dependency DAGs with cycle detection, and lock-free single-flight.
- Engineered the queue's **crash-isolation** so no job hangs non-terminal: each job runs in a spawned (not forked) supervisor that survives a segfault or OOM-kill and reaps the child's process group, while a heartbeat sweeper force-fails jobs whose supervisor dies or freezes (catching a frozen-but-alive one a PID check would miss). Chaos tests `SIGKILL` live supervisors and kill the database mid-job to prove it.
- Built **GPU-memory admission control** so concurrent LLM jobs can't overcommit one card: a per-pool Postgres advisory lock admits a job only when its declared context-token cost fits the pool's live budget, and a runtime `READ COMMITTED` requirement makes a claimer that waited on the lock re-price against just-committed usage instead of a stale snapshot.
- Built the in-house **agent runtime** the platform's security agents run on (LangChain/LangGraph): context compaction that counts real tokens and fails closed instead of overflowing, plus tool-output eviction, loop detection, and per-run container sandboxing (fresh, network-isolated, self-reaping).
- Built **PurpleHaze**, an autonomous penetration-testing product: a local LLM drives Metasploit and real offensive tooling through a full multi-step engagement, holds session state across long runs, and produces compliance-grade OWASP and NIST reports, with no human in the loop.
- Sole owner of AtlasCyber's AWS**, end to end across GovCloud and commercial partitions: EC2 (with a reverse-proxy/tunnel fleet), S3, IAM, VPC networking, and CloudFront.
- Rewrote the **network-capture sensor** from scratch in Rust, memory-safe workspace-wide (`#![forbid(unsafe_code)]`): a headless Linux daemon around a crash-safe SQLite-WAL upload queue that persists each capture before upload and survives a restart without dropping data, replacing the prior lossy in-memory queue.

CrunchAtlas — Software Engineer, AtlasConnect | Mar 2025 – Nov 2025

Sole developer and maintainer of an inherited codebase — Django REST, React/Vite, self-hosted LLM stack.

- Built the **document-ingestion pipeline** for deal intake: multi-format extraction (PDF/DOCX/PPTX/DocSend) with selective parallel OCR on scanned pages, feeding concurrent LangChain chains under a managed token budget.

PROJECTS

gloss — Python, SQLite/FTS5, Ollama

- Built a hybrid retrieval engine that fuses BM25 with a local-embedding vector channel via reciprocal-rank fusion to return a source document's verbatim, cited passages, tuned against a 266-case evaluation gated by a paired significance test (hit@5 0.94, hit@1 0.71).

Celery Fork-Safety Investigation — Python, Celery, GDB, LLDB, Docker

- Root-caused a Celery prefork fork-safety bug to the register level: a Rust-backed vector store initialized before `fork()` left the child inheriting a lock whose owner didn't survive the fork — a permanent futex deadlock on Linux (traced in GDB) and a libdispatch semaphore crash on macOS (traced in LLDB). One cause, two OS-level failure modes.

HONORS

- 1st Place / Best Hack — Bitcamp 2024**, UMD (largest US East Coast hackathon) — TermBook
- Best Educational Hack & 2nd Overall — HackUMBC 2024** — Testif.AI

SKILLS

- Backend & Infrastructure:** Django, FastAPI, PostgreSQL, MongoDB, SQLite, Redis, Celery, Linux, AWS
- DevOps & CI/CD:** Docker (multi-stage, BuildKit), GitHub Actions, Jenkins, supervisord, systemd
- Languages:** Python, Rust, TypeScript, SQL, Bash
- AI & LLM:** LangGraph/LangChain, llama.cpp, llama-swap, Ollama, LLM serving, GPU memory management

EDUCATION

Salisbury University — B.S. Computer Science, Minor in Mathematics | Aug 2021 – May 2025