DYLAN HAMPTON

St. Peters, MO · dylanjosephhampton@gmail.com · linkedin.com/in/dylanjosephhampton · dylan.markset.io
Open to remote or relocation · U.S. citizen, no sponsorship needed · eligible to obtain a Secret clearance

SUMMARY

Software engineer who builds LLM platforms in regulated cloud environments and the AI workflows business teams use every day. At a defense contractor, built the team's LLM gateway on Amazon Bedrock with identity and audit built in, and message encode/decode and translation logic between military data standards. Through my consultancy, shipped Claude skills to a contractor's project managers and tested AI output against real completed work.

SKILLS

- **AI engineering:** LLM APIs and tool use, Claude skills and hooks, agents, LiteLLM gateways, Amazon Bedrock and Guardrails, evals against real completed work, directing Claude Code and Codex and reviewing their output, local inference (llama.cpp)
- **Safety and security:** ITAR/CUI environments, FedRAMP High, DoD IL4/IL5, identity from authentication, spending limits in code, human approval before side effects, audit logging
- **Integration:** standards-based message encode/decode and translation, conformance validation, REST APIs (FastAPI, Java HttpServer), connectors (SharePoint, Outlook)
- **Languages and tools:** Python, TypeScript/JavaScript, Java, C++, C#, Bash, PowerShell, AWS (GovCloud, Bedrock, IAM, CloudWatch, Athena), Docker, Git, GitLab, Linux, SQLite, MongoDB

EXPERIENCE

Software Engineer, Stauder Technologies St. Peters, MO · May 2024 – Present

AI platform (ITAR, IL4/IL5 AWS GovCloud)

- Built a LiteLLM gateway in front of Amazon Bedrock with a companion service that stamps every request with the user and project, so usage shows up in AWS's own logs. Identity comes from the authenticated key, not headers a client could fake, and only metadata is attached, never prompt content.
- Found why the obvious approach failed (request-body metadata is stripped on one API route) and fixed it with a pre-call hook.
- Designed and directed a FastAPI bridge (AI-assisted implementation) routing Codex CLI traffic through the gateway to Bedrock. Rolled out Codex CLI on Windows, fixing a PowerShell 5.1 incompatibility. Took part in model selection and migration planning for a multi-provider gateway.
- Set up Bedrock Guardrails and usage policies supporting FedRAMP High. Built the onboarding path (bootstrap scripts, data-handling rules, a shared prompt library) and ran workshops for ~30 engineers; the whole team now uses it every week.

Data integration between standards (C++, Java, C#, TypeScript)

- Built message encode/decode for VMF (MIL-STD-6017 and earlier revisions), Link 16 (MIL-STD-6016) and JREAP-C (MIL-STD-3011) across two client applications, and gateway logic translating messages between standards and revisions, including Cursor on Target. Validated conformance against each standard.

Founder, MarkSet (AI workflow consultancy) St. Louis, MO · Apr 2026 – Present

- Ran a paid hands-on Claude workshop for a mechanical contractor's project managers and estimators on their live jobs, then shipped three Claude skills to their Enterprise org: a job briefing that pulls each job's files and mail into one page, a morning report across active jobs and open bids, and a writing-voice profile.
- A dry run with a real PM caught the briefing stating things that weren't true, so every line now carries an evidence tag pointing to its source. Found the SharePoint connector silently stops at about row 1,950, so the report reads a local copy with a row checkpoint.
- Tested an estimating agent against completed bids: within 0.3% of the award on one job, a 36.6% miss on another, traced to missing calibration for that job type.
- Built a skill that turns a prospect's company URL into a fit analysis, workflow maps and a working demo dashboard. Wrote a client data guide for customers that restrict AI use: where data goes, what the provider keeps, and a per-job AI rule.

PROJECTS

- **Prediction-market agent.** Scans every open Kalshi market (5,000+) on a schedule; 160+ tests against a simulated exchange; running a paper month before any real money. Spending limits live in code, and judgment calls go to a human approval inbox that records each decision's reason.
- **LearnFrame.** A repo-level agent setup I direct daily: context files, 12 shared skills, hooks, and a nightly self-improvement job that has run since June 2026 (430+ commits, built with AI coding agents). It has set up and drives ten separate projects.
- **On-device phone agent.** llama.cpp serving Gemma 4 E2B on a Galaxy S25 Ultra through Termux at a measured ~16 tokens/sec, with an allowlisted shell and a typed "yes" required before any side effect.

EDUCATION

B.S., Computer Science, Missouri University of Science and Technology May 2024