

HAMZAH ASSAF

Software Developer | Full-Stack | Backend | Cloud

Beirut, Lebanon • hamzahassaf08@gmail.com • +961 3741213 • linkedin.com/in/hamzah-assaf • hamza2497.github.io

PROFILE

Full-stack developer with hands-on LLM evaluation experience, scoring model outputs daily at Invisible Technologies and shipping AI tooling built from that work. Shipped PIP (a Claude plugin for AI-assisted learning) and an MCP eval server exposing five evaluation tools, both production-deployed. .NET / Python / React, AWS certified.

PROJECTS

PIP *Repo-Aware Learning Layer for AI-Assisted Development - Claude Plugin* 2026

- Solves a real blind spot in AI-assisted development: AI coding tools make you ship faster but teach you nothing. PIP reads your actual repo, generates a concept tree from your code and goals, then teaches each concept through conversational checkpoint loops with AI-evaluated understanding checks, deduplicated across projects.
- Designed a token-saving subagent architecture that delegates resource-heavy codebase parsing to lightweight models, reserving frontier-model tokens for complex reasoning, with zero external infrastructure and state kept locally in the repo.
- Plugin authored as composable skills over the host's native file and git tools (Python). A deployed FastAPI + React web app serves as a live demo of a fraction of the plugin's capability. [Live Demo](#) | [GitHub](#)

MCP Eval Server *AI Evaluation Tool - MCP Plugin + Web App* 2026

- Built an LLM output evaluation server exposing five tools (rubric scoring, factual accuracy, relevance, logical consistency, and head-to-head comparison), callable directly from Claude Desktop as an MCP plugin.
- Delivered two interfaces: a FastAPI + Jinja2 web app (Gemini as judge, live on Render) and a local Claude Desktop plugin (Claude as judge, no API key required); backed by 19 passing tests.
- Built in Python using FastMCP/Pydantic, drawing on hands-on LLM evaluation work. [Web App](#) | [GitHub](#)

Library Management System *Full-Stack - Angular · .NET · JWT · Docker · CI/CD* 2025

- Full-stack app with an Angular frontend and a .NET 8 / ASP.NET Core REST API: JWT auth, Google SSO, role-based authorisation, EF Core + PostgreSQL, Docker deployment on Render, GitHub Actions CI/CD pipeline. [GitHub](#)

TECHNICAL SKILLS

Backend: C#, .NET 8, ASP.NET Core, Python, FastAPI, REST APIs, Entity Framework Core, Dependency Injection, Repository Pattern, Async/Await, Middleware, LINQ, Multithreading

AI & APIs: LLM output evaluation, MCP server development, FastMCP, Google Gemini API, Pydantic, Swagger / OpenAPI

Frontend: React, Angular, TypeScript, JavaScript, HTML, CSS, Tailwind CSS, Vite

Databases: PostgreSQL, SQLite, SQL Server, EF Core Migrations

Cloud & DevOps: AWS (EC2, S3, IAM, VPC, CloudWatch, CloudFormation, CLI), Docker, GitHub Actions, CI/CD, Render, Vercel

Auth & Security: JWT Bearer Tokens, Google SSO / OAuth, Role-Based Authorization

Testing: xUnit, pytest, Playwright

Tools: Git, GitHub, AI-assisted development (Claude Code)

EDUCATION & CERTIFICATIONS

- B.S. in Computer Science, Lebanese International University, Beirut (2021-2024)
- AWS Certified Cloud Practitioner: EC2, S3, IAM, VPC, CloudWatch, CloudFormation, CLI (Dec 2024)
- Docker Certified Associate (Aug 2024)

EXPERIENCE

AI Quality Engineer *Invisible Technologies (Contractor)* 2025 - Present

- Scored LLM outputs across instruction-following, factual accuracy, logical consistency, and edge-case failure detection, and audited training data using structured rubrics, applying consistent quality standards across diverse task types.
- Served as quality gatekeeper for the evaluation team, reviewing and approving other trainers' work and setting the accuracy benchmark across the team.

Technical Support Manager *Securealm* 2018 - 2019

- Diagnosed and resolved customer-reported technical issues using remote support tools and CRM; documented detailed reproduction steps and escalated to engineering with full context for resolution.

Languages: English (fluent) • Arabic (native)