Zain Dana Harper

Kent, Washington | zaindharper@gmail.com | harperz9.github.io
linkedin.com/in/zaindanaharper | github.com/HarperZ9 | harperz9.github.io/published-work.html

Curriculum Vitae | Security Evaluation, AI Evaluation, Graphics, and Developer Tools

Profile

Self-directed software builder who writes tools that test AI systems and the software around them, and graphics work that people install. Elder ENB, a lighting preset for Skyrim Special Edition, has 996,649 downloads on Nexus Mods. About 13 years of personal programming. Career transition began June 2026 after eleven years of full-time production arboriculture; no professional programming experience.

Seeking red-team, purple-team and security-evaluation work, AI evaluation roles, and graphics engineering. Willing to relocate to London; available for regular travel to San Francisco.

Selected Evidence

- Elder ENB: 996,649 total downloads, 176,359 unique downloads and 2,413 endorsements on Nexus Mods, read October 23, 2026. First uploaded April 21, 2024; last updated February 20, 2026, at version final-elder-9. nexusmods.com/skyrimspecialedition/mods/117327
- Open source: 37 merged engineering pull requests in 30 external repositories, with 17 closed without merge and 14 open, counted from the GitHub API on October 23, 2026. harperz9.github.io/career/open-source-census.json

Flagship Projects

- Flywheel runs AI tasks with the model you choose and checks each answer against a task contract. In 1.9.0, behavior-check flagged 20 of 20 planted behavior changes and 0 of 50 unchanged pairs. github.com/HarperZ9/flywheel/releases/tag/v1.9.0
- CMMNS (formerly Daylight Commons) is a browser extension that shows the sources behind replies on nine social platforms. Its slur sign has 91.0% precision, 3,995 of 4,392, on public human-labelled corpora. github.com/HarperZ9/cmmns/blob/main/project-docs/EVALUATION.md
- Articulate is prose analysis and model-assisted editing for the command line, editors and CI. In 0.14.0, Markdown masks removed 5,123 findings on table separator rows across 3,000 files and added none. github.com/HarperZ9/articulate/releases/tag/v0.14.0
- BuildLang is a systems language in which a program declares what it may touch, and the compiler checks it. Version 1.10.0 is on crates.io. crates.io/crates/buildlang
- raw-native is a C++ renderer with a CPU reference and D3D12 and WebGPU backends. In release 0.8.0, default renders were byte-identical across the Windows, D3D12, Linux and WebAssembly builds. github.com/HarperZ9/raw-native/releases/tag/v0.8.0
- Low Water is an original narrative game in development on raw-native, in a private repository.

Self-Directed Projects

About 13 years of self-directed programming | Career transition began June 2026

Sole proprietor since September 2026; no paid software clients

- Own the architecture, implementation, integration, tests, documentation, and release work for Flywheel and related tools. Work spans Python services and command-line tools, Flutter desktop software, Rust systems tooling, and browser interfaces.
- Develop model-agnostic workflows, evaluation harnesses, and controlled-action interfaces whose records can be checked without a model approving its own answer.

Personal programming began with game modding, reverse engineering, bot and rotation logic, and runtime-memory tooling, then developed into systems engineering, graphics, compilers, and AI infrastructure.

Freelance Technical Writer, Documentation, and Product Operations

Independent practice | Began 2017; project-based work

- Produced technical and compliance documentation, API and implementation guides, proposals, release notes, and operational playbooks. Turned implementation details into material usable by customers, developers, and reviewers.

Employment

Commercial Arboriculture and Field Operations | Legendary Tree

Seattle area | April 25, 2015 to June 2, 2026

- Combined ground-based tree-care work and safety coordination with client intake, site assessments, estimates, proposals, scheduling, and crew/vendor communication.
- Monitored rigging systems, clearances, and changing conditions while supporting climbers from the ground. Explained work options and constraints to clients and maintained continuity between the proposed scope and field execution.

Technical Networking Support | Xbox/Microsoft contract

Subcontracted through Stream/Convergys | Wilsonville, Oregon | 2014 to 2015

- Diagnosed connectivity and account issues by phone and chat using TCP/IP, DNS, NAT, firewall, and router troubleshooting; documented repeatable resolutions and customer-facing instructions.

Security Evaluation: Red Team and Purple Team

No paid red-team or purple-team engagement to date. This work runs on my own systems and test targets, with planted faults and attacks.

Planted controls in release checks

Released a Flywheel monitor-admission check over 2,676 fixed test items with six planted controls and a second, independent math kernel. github.com/HarperZ9/flywheel/releases/tag/v1.5.0

Accountable Surface | Governed actions and computer-use adapters

Built operator-granted action interfaces for filesystem, command, API, browser, and Windows UI Automation work. Actions carry explicit scope and postconditions, journals preserve decisions and observed effects, and fake drivers make browser and desktop scenarios testable without live actions.

On PyPI as accountable-surface 0.3.1; alpha. pypi.org/project/accountable-surface

Phantom | Hardware-identity inspection and reversible system changes

Developed a Rust utility for Windows and Linux with read-only audit, backup-before-write, validation, restoration, and platform packages. Release v1.1.1 changes and restores user-level identifiers. github.com/HarperZ9/phantom/releases/tag/v1.1.1

Controlled assessment and detection engineering

Private systems work includes Array for approval-gated assessment orchestration, Lab/Seed for controlled-lab fixtures and detection-engineering workflows, and ORCA for engagement state, findings and reports. Source and operational details are private.

AI Evaluation

Flywheel | Model-agnostic AI workstation and coding harness

Primary owned platform. Flywheel 1.9.0, released October 15, 2026, is on PyPI as flywheel-verify, installs through Homebrew or Scoop, and ships a Windows installer with a bundled engine.

- Routes tasks to hosted or local models through an OpenAI-compatible gateway, with command-line, desktop, and API entry points.
- Provides Rowan, a desktop assistant that turns requests into recorded tasks, with permission-checked coding and tool-use workflows and session state.

Terminal State Fixtures and Crucible
Test infrastructure for evaluating AI agents: 324 terminal-state fixtures and 337 component cases that tell a failed answer apart from a provider failure or missing evidence. Ran a pinned METR count_odds task through the Inspect-to-Flywheel bridge with three deterministic controls. github.com/HarperZ9/terminal-state-fixtures

CMMNS (formerly Daylight Commons) | Browser extension for checking where replies come from
A browser extension that shows the sources behind replies on X, YouTube, Reddit, Facebook, Instagram, Bluesky, Threads, TikTok and Mastodon. Runs locally with no account. Open source (AGPL-3.0). Version 3.2.2, released October 23, 2026. github.com/HarperZ9/cmmns/releases/tag/v3.2.2

Articulate | Prose analysis and editing
Prose analysis and model-assisted editing for the command line, editors and CI. It names each writing pattern it finds and the line it is on, in the terminal, in an editor through its language server, or as SARIF in code review. Version 0.16.0 is on PyPI as articulate-writing. pypi.org/project/articulate-writing

Graphics and Rendering

Elder ENB | Lighting and post-processing preset for Skyrim Special Edition
An ENBSeries preset with a layered design: weather presets set the atmosphere, and stackable overlays adjust style, lighting, effects and post-processing. It ships 38 atmosphere presets for the KreatE weather system, updated from April 2024 to February 2026. ENBSeries and the required mods belong to their authors. nexusmods.com/skyrimspecialedition/mods/117327

raw-native | C++ renderer with a CPU reference and checked GPU backends
Release 0.8.0, October 14, 2026, renders on a CPU reference and through D3D12 and WebGPU, and passed 60 of 60 tests with MSVC and GCC. github.com/HarperZ9/raw-native/releases/tag/v0.8.0

Merged since 0.8.0 and not yet released: a Vulkan 1.3 backend whose 600 presented frames on lavapipe match the offscreen and D3D12 WARP renders hash for hash; the painterly shader ported to the C++ core, equal to the JavaScript version across fifteen cases; a tilt-shift lens that repaints out-of-focus areas with coarser brushwork and no blur, about 0.86 ms a frame at 1080p on an RTX 4090; and a character asset toolchain that turns brush strokes into signed distance fields, game meshes, UV atlases and baked maps.

Low Water | Original narrative game in development
An original narrative game written and built on my own engine, with a deterministic combat simulation. A C++ port onto raw-native is in progress in open raw-native pull requests. The repository is private.

Earlier graphics and color work
Earlier HLSL and C++ modding covered post-processing, Direct3D 11 integration, shader debugging and GPU traces, with public projects SkyrimBridge, ENB Runtime Core, Brender Archival and Engine Revival. Color tools include Build Color, Calibrate Pro and Studio Engine.

Developer Tools and Languages

BuildLang | Rust compiler and developer tooling
Developed a systems-language toolchain with typed capability effects, type inference, parser and intermediate-representation work, native C output, shader output, and language-server support. Version 1.10.0 is on crates.io, published October 23, 2026; the 1.8.0 release that day added Box<T> and recursive types. crates.io/crates/buildlang

EMET | Byte-integrity witnessing
Developed Python, Rust, Node, and Go implementations that check whether bytes still match their source, sharing one set of conformance cases. On PyPI as emet. pypi.org/project/emet

Research intake and workspace tooling

Built Gather for research intake from web, scholarly and document sources, Index for workspace and symbol maps, and Forum for running teams of agents. On October 23, 2026, Gather 2.4.0, Crucible 1.6.0, and EMET 1.3.1 shipped to PyPI and Canon 0.7.0 to GitHub. Companion tools include Relay for coding agents, Plexus for tool wiring, Mneme for agent memory, Canon for assistant instructions, and Bulletin, a message board for agents.

Selected Open-Source Contributions

Six of 37 merged engineering pull requests in 30 external repositories. Full list: harperz9.github.io/career/open-source-census.json

- DeepEval #2822: fixed unhashable tool-output and nested-parameter failures in tool-correctness evaluation; merged July 2, 2026. github.com/confident-ai/deepeval/pull/2822
- TOMLKit #549: preserved dotted-key siblings when replacing a value with an array of tables; merged July 14, 2026. github.com/python-poetry/tomlkit/pull/549
- Datasette #2815: corrected malformed composite-key requests from HTTP 500 or incorrect HTTP 200 to HTTP 400; tested too few and too many key components in HTML and JSON. Merged July 7, 2026. github.com/simonw/datasette/pull/2815
- Free Law Project #820: documented runnable database-free fast tests; merged August 25, 2026. github.com/freelawproject/litigant-portal/pull/820
- Hebbian Robotics #157: tested reserved-column drift against the actual queryable episode schema; merged August 25, 2026. github.com/Hebbian-Robotics/hflow/pull/157
- Voxwire #44: added offline regression coverage for a six-event degraded voice turn, matching saved PCM audio and JSONL records to emitted events. Merged August 15, 2026. github.com/vectorvoyager358/voxwire/pull/44

Technical Competencies

Languages and runtimes: Python, Rust, C++, Dart/Flutter, TypeScript/JavaScript, Lua, HLSL, C#, PowerShell, and Bash. Additional cross-language conformance work includes Go and Node implementations.

Security evaluation: planted-fault and control design, approval-gated assessment workflows, detection engineering fixtures, vulnerability management and policy audit, and evidence handling for findings.

AI infrastructure and quality: multi-provider/local-model routing, tool-use loops, task decomposition, MCP, evaluation design, deterministic fixtures, regression testing, worker/verifier separation, failure analysis, and terminal-state scoring.

Controlled actions: capability effects, approval grants, hash-chained journals, offline replay, and conformance suites.

Software delivery: Git/GitHub, GitHub Actions, Python packaging, command-line design, Windows/Linux installers, CMake, test automation, and release checksums.

Graphics and domain knowledge: D3D11/D3D12, Vulkan, HLSL, WebGPU/WGSL, ENBSeries presets, painterly and stylised post-processing, signed distance fields, CPU reference rendering and GPU parity checks, color-space and tone-mapping work, GPU instrumentation, and applied arboriculture/horticulture from field practice.

Working Practice

AI-assisted development is part of the workflow. Architecture, integration, source review, test design, release decisions, and responsibility for the submitted work remain mine.

Research Interests

Independent verification of AI evaluations, drift in live systems, records of agent actions, human approval for consequential actions, and robotics.

AI Accountability, Disclosure, and Public-Interest Writing

Independent essays on AI accountability, disclosure, conflicts of interest, and human consequences, relevant to public-interest and whistleblower-protection work.

- Who Knew First, September 2026. Directed and published an AI-assisted investigation comparing disclosure timelines, evidence access, institutional incentives, and evaluator conflicts across nine reported AI-agent incidents. harperz9.github.io/who-knew-first.html
- An Open Letter on Checking the Machines, September 2026. Proposed contestable AI decisions, repeatable evidence checks, and accountable responses to objections. harperz9.github.io/checking-the-machines.html
- Borrowed Ground, September 2026. Philosophical essay connecting human interdependence to responsibility, privacy, participation, and the limits of technical evidence as authority over people. Substantial AI assistance disclosed. harperz9.github.io/borrowed-ground.html
- The Sandbox Was Never Just a Box, September 2026. Compared public security reports to examine shared-service exposure and the independence of logs, monitors, and scorers in AI-agent evaluation. harperz9.github.io/the-sandbox-was-never-just-a-box.html

Selected Publications and Research Records

Independent systems papers, preprints, research notes, philosophy papers, a working paper, and archived manuscripts with DOI records. Not peer reviewed.

- EMET: An Authority-Incapable Byte-Level Integrity Witness. Systems paper, 2026. Byte-integrity witnessing and the separation of observation from authority. doi.org/10.5281/zenodo.21230267
- BuildLang: Accountable Compute via Typed Capability Effects. Systems paper, 2026. Capability effects and accountable compilation. doi.org/10.5281/zenodo.21231253
- Witnessed Independence: Recording Whether a Verifier Graded Its Own Work. Preprint, 2026. doi.org/10.5281/zenodo.21232206
- Proof Packets: A Derive-Don't-Trust Envelope for Accountable Agent Actions. Preprint, 2026. doi.org/10.5281/zenodo.21231406
- The Personhood-Gate Handoff. Research note, 2026. doi.org/10.5281/zenodo.21234475
- Re-Perceived Effects. Research note, 2026. doi.org/10.5281/zenodo.21231311
- Faithfulness to a Named Criterion Is the Conserved Quantity Across Lossy Substrates. Research note, 2026. doi.org/10.5281/zenodo.22768398
- The Witnessing Spine. Archived research corpus, 2026. doi.org/10.5281/zenodo.20778927
- Conferred Existence. Long-form philosophical manuscript and archived corpus, 2026. doi.org/10.5281/zenodo.20773724
- Sourcehood Without a Source: The Arity Gap and the Free Will Debate. Philosophy paper, edition of 3 October 2026. doi.org/10.5281/zenodo.23126117
- The Seam That Holds While It Is Crossed. Philosophy paper, edition of 3 October 2026. doi.org/10.5281/zenodo.23126357
- Self-Given Is Not Self-Grounding. Philosophy paper, edition of 3 October 2026. doi.org/10.5281/zenodo.23126458
- The Conservation of Faithfulness. Working paper, edition of 3 October 2026. doi.org/10.5281/zenodo.23126484

Education

High School Diploma | Wilsonville High School, Wilsonville, Oregon | 2013

Self-directed study in systems programming, compilers, graphics, color science, AI evaluation, epistemics, and philosophy of mind. Practical arboriculture and plant-science knowledge comes from field work.

ORCID: orcid.org/0009-0001-7175-5393 | harperz9.github.io/publications.html | harperz9.github.io/portfolio.html

Credentials and Continuing Education

2026 completions. Verification: harperz9.github.io/career/source-ledger.md.

Qualys Certified Specialist: Policy Audit; Custom Assessment and Remediation; Vulnerability Management Detection and Response; Qualys Kubernetes and Container Security; Qualys TotalCloud Exam; Endpoint Detection and Response | Issued 2026-09-30; expire 2028-09-29

Microsoft Applied Skills | Earned September 2026

- Secure AI solutions in the cloud
- Get started with identities and access using Microsoft Entra
- Secure storage for Azure Files and Azure Blob Storage
- Implement security through a pipeline using Azure DevOps | Intermediate; 95% assessment

Advanced Microsoft Learn Coursework

SC-100: Design security solutions for applications and data. Four module assessments and the case study.

Design solutions that align with MCRA and MCSB. Advanced Microsoft Learn module completed September 30, 2026: security architecture, Zero Trust, insider and supply-chain risk, and AI security controls.

Linux Foundation Course-Completion Certificates

LFEL1012: Secure AI/ML-Driven Software Development; LFS120: Conversational AI: Ensuring Compliance and Mitigating Risks; LFD121: Developing Secure Software

LFS182: Securing Your Software Supply Chain with Sigstore; LFEL1001: Understanding the EU Cyber Resilience Act (CRA); LFS118: Ethical Principles for Conversational AI

LFEL1006: Securing Projects with OpenSSF Scorecard; LFEL1005: Security Self-Assessments for Open Source Projects; LFEL1007: Automating Supply Chain Security: SBOMs and Signatures

LFD125: Security for Software Development Managers; LFS180: Introduction to DevSecOps for Managers; SKF100: OWASP Top 10 Security Threats

LFS184: Introduction to JavaScript Security; LFC108: Cybersecurity Essentials; LFEL1004: Authentication and Authorization for Web/API

LFEL1008: Migrating AI Workflows Across Platforms; LFS147: Introduction to AI/ML Toolkits with Kubeflow

Additional Coursework and Skill Badges

Cisco Networking Academy: Ethical Hacker | Course certificate and learning badge, 2026-09-30; 100% knowledge exam.

Google skill badges, completed 2026-09-30: Google DeepMind: Train A Small Language Model (Advanced; 100/100, four coding checkpoints); Use Agent Skills with Multi-Agent Systems (Intermediate course/Advanced challenge; 100/100, six checkpoints). Multi-agent lab: five Cloud Run services deployed; model-access policy blocked local and cloud generation.

Google BigQuery skill badges, issued 2026-10-01: Perform Predictive Data Analysis in BigQuery (100/100, seven checkpoints; Advanced challenge); Create ML Models with BigQuery ML (100/100, four checkpoints; Intermediate challenge).

Google course-completion badges, completed 2026-09-30: Machine Learning Operations (MLOps) with Agent Platform: Model Evaluation (Intermediate); Google Security Operations - Deep Dive (Advanced; 96% knowledge assessment).

Anthropic Claude Academy completion badges: Model Context Protocol: Advanced topics; Claude Code in action; Building with the Claude API; Claude with Amazon Bedrock | Issued September 30, 2026

UNESCO: Global MOOC on the Ethics of AI | Intermediate course certificate, 2026-09-30; 100%.

BlueDot Impact: AGI Strategy | Course completed in 2026

Developed an AGI Strategy action plan for label-independent AI incident notice, independently reviewable evidence, accountable escalation, and tests for misleading verification results.

Coursework: MCRA/MCSB, Zero Trust, OWASP web Top 10, CRA. Research references: NIST AI RMF/GenAI Profile, ISO/IEC 42001/23894, ISO 37002. Scope: harperz9.github.io/career/standards-reference.md.