

Zain Dana Harper

Portfolio. The engineering, the upstream ledger, and the writing and analysis samples.

Seattle, Washington | zaindharper@gmail.com | harperz9.github.io | github.com/HarperZ9 (<https://github.com/HarperZ9>)

How to read this index

A working index of writing and analysis samples, plus the engineering work that makes up the flagship line. Every claim links to a public artifact you can open: a crate, a marketplace listing, a repo, a release, or a live page. Maturity is stated plainly as beta, current tagged release, versioned public package, or source version, never inflated.

Material that could not be made public-safe was left out on purpose. See What is not here at the end.

Fourteen flagship engines

Fourteen public, inspectable engines share one Flywheel thesis: each stands on its own, shows its work, and can pass verified outputs forward. Full detail and recorded workflows are on the flagships page.

- **flywheel** · route and verify any model. Local or hosted, it checks what it can and returns a re-checkable receipt. Public, FSL-1.1-MIT. Open.
- **telos** · the engine. A person and a model perceive the same thing, shape it, and verify every step. Open the Studio.
- **index** · the map. Repository-inventory and dependency-graph engine for many-repo workspaces. index-graph 2.9.0; Development Status Beta. Open.
- **gather** · the intake. Accountable research intake with provenance receipts and digest seals. gather-engine 1.6.1, the current tagged release in the local public clone. Open.
- **forum** · the orchestrator. Runs a team of agents and records every step in a witnessed causal ledger. forum-engine 1.13.0, a versioned public package. Open.
- **crucible** · the judgment. Registers a thesis, steelmans its claims, measures them against a substrate, and refines the weakest axis. crucible-bench 1.2.0, a versioned public package. Open.
- **emet** · byte integrity. Re-derives a file's bytes from scratch without altering it. EMET 1.1.0, four independent implementations, public on PyPI. Open.
- **buildlang** · the compiler. A typed-effects language: functions declare the effects they may perform and the lifetimes of the references they return, checked and lowered through a verified C path. v1.2.0 manifest source version; additional backends remain experimental. Open.
- **learn** · the learning aid. Turns supplied material into a course with retrieval practice and witnessed grading. Learn 1.6.0, zero-dependency Node. Open.
- **relay** · code with any model. A zero-dependency coding agent across local, subscription, and API endpoints with a gated tool loop. Relay 0.1.0, public source. Open (<https://github.com/HarperZ9/relay>).
- **plexus** · wire tools together. Discovers tool capabilities and assembles inspectable producer-to-consumer pipelines. Plexus 0.2.0, public source. Open (<https://github.com/HarperZ9/plexus>).
- **mneme** · agent memory. Layered memory and hybrid retrieval with provenance carried through recall. Mneme 0.1.0, public source. Open (<https://github.com/HarperZ9/mneme>).
- **studio-engine** · generate worlds. Seeded shaders, sound, and motion timelines that redraw from the same inputs. Studio Engine 0.2.0, public source. Open (<https://github.com/HarperZ9/studio-engine>).
- **build color** · measure color. A bounded color-science workbench for perceptual spaces, color difference, HDR tone mapping, gamut tools, and ICC profiles. Build Color 1.0.2, PyPI, beta. Open.

Private-line platform: Checkpoint, Vault, Runtime, and Boundary (<https://github.com/HarperZ9/behavior-transform.io>) are public where safe; Lab (Seed) and Ledger (Sofer) remain private/proprietary until their public-safe split is complete. The line carries release-checkpoint receipts, MCP and CLI contracts, package checks, and public-safe documentation.

Contributions to projects I do not own

Work sent into other people's codebases, where someone else sets the standard and decides. Every line below is one pull request you can open: the diff, the review, and the outcome. Merged means a maintainer accepted that change. It does not mean the project uses, endorses, or has any relationship with the rest of this portfolio.

60 pull requests

25 merged

22 open

13 closed without merge

Every pull request authored by HarperZ9 in a repository not owned by HarperZ9. The counts are read straight from the GitHub API by tools/pr_census.py and a test fails the build if this page disagrees with the live record, so nothing here is transcribed by hand. **14** of them only add a row to a public link list. Those are held apart at the end of this section, because adding a row to an index is not a contribution to a project. That leaves **46** code and documentation pull requests: **23** merged, **12** open, **11** closed without merge.

Merged

23 PULL REQUESTS · 20 REPOSITORIES

Defect repair, missing validation, and test coverage, accepted by the maintainer.

- **simonw/datasette** · return 400 instead of 500 for wrong-arity composite primary-key row URLs. **#2815** (<https://github.com/simonw/datasette/pull/2815>)
- **python-poetry/tomlkit** · an array of tables replacing a dotted key no longer swallows the next sibling key. **#549** (<https://github.com/python-poetry/tomlkit/pull/549>)
- **pydantic/pydantic-ai** · preserve `FileUrl.force_download` through UI round-trips. **#6205** (<https://github.com/pydantic/pydantic-ai/pull/6205>)
- **confident-ai/deepeval** · stop the tool-correctness metric crashing on unhashable tool outputs. **#2822** (<https://github.com/confident-ai/deepeval/pull/2822>)
- **dgilland/pydash** · `to_list()` returns a list for dicts, not a `dict_values` view. **#250** (<https://github.com/dgilland/pydash/pull/250>)
- **python-grimp/grimp** · raise `ModuleNotPresent` for absent modules in direct-import queries. **#303** (<https://github.com/python-grimp/grimp/pull/303>)
- **python-grimp/grimp** · raise `ValueError` when one module is assigned to more than one layer. **#304** (<https://github.com/python-grimp/grimp/pull/304>)
- **python-grimp/grimp** · clear the ruff and clippy failures that were holding the main branch red. **#305** (<https://github.com/python-grimp/grimp/pull/305>)
- **open-multi-agent/open-multi-agent** · conformance fixture for mid-task recovery. **#452** (<https://github.com/open-multi-agent/open-multi-agent/pull/452>)
- **tylerjrbuell/reactive-agents-ts** · seeded end-to-end determinism coverage for replay. **#196** (<https://github.com/tylerjrbuell/reactive-agents-ts/pull/196>)
- **tylerjrbuell/reactive-agents-ts** · return replayed steps through the harness handle. **#197** (<https://github.com/tylerjrbuell/reactive-agents-ts/pull/197>)
- **redxzeta/runtrail** · preserve bridge startup diagnostics instead of discarding them. **#154** (<https://github.com/redxzeta/runtrail/pull/154>)
- **sjh9714/mergewarden** · replay fixtures for reusable-workflow pinning. **#107** (<https://github.com/sjh9714/mergewarden/pull/107>)
- **adammmmmm/LoopCompass** · an evaluation benchmark suite. **#19** (<https://github.com/adammmmmm/LoopCompass/pull/19>)
- **AjnasNB/maqam** · offline coverage for Google ADK function-tool routing. **#28** (<https://github.com/AjnasNB/maqam/pull/28>)
- **Liuyanfeng1234/governance-spec** · duplicate-key rejection fixtures. **#1** (<https://github.com/Liuyanfeng1234/governance-spec/pull/1>)
- **0xZR/freellmpool** · a capacity-status edge fixture. **#79** (<https://github.com/0xZR/freellmpool/pull/79>)
- **michalstrnadel/scorekeeper** · a report output and a Windows file-lock fallback. **#7** (<https://github.com/michalstrnadel/scorekeeper/pull/7>)
- **morganlinton/Albatross** · support external eval fixture files. **#13** (<https://github.com/morganlinton/Albatross/pull/13>)
- **AvAdiii/rewardspy** · an ASCII fallback for dashboard charts. **#7** (<https://github.com/AvAdiii/rewardspy/pull/7>)
- **sguruvar/ai-inference-observability-eks** · fix the local MCP endpoint. **#1** (<https://github.com/sguruvar/ai-inference-observability-eks/pull/1>)
- **lliangcol/llm-plugins-fusion** · a worked sample workflow-evaluation record for the docs. **#104** (<https://github.com/lliangcol/llm-plugins-fusion/pull/104>)
- **haeliotang/agent-runtime-observatory** · content-only golden traces, so re-recording leaves a clean diff. **#42** (<https://github.com/haeliotang/agent-runtime-observatory/pull/42>)

Open

12 PULL REQUESTS · 12 REPOSITORIES

Submitted and awaiting a maintainer. Open is not accepted, and none of these are counted as merged.

- **simonw/datasette** · give each non-blocking write a distinct task id. #2861 (<https://github.com/simonw/datasette/pull/2861>)
- **simonw/llm** · timeout options for the OpenAI Responses client. #1506 (<https://github.com/simonw/llm/pull/1506>)
- **drizzle-team/drizzle-orm** · preserve commas inside MySQL enum values in snapshots. #5961 (<https://github.com/drizzle-team/drizzle-orm/pull/5961>)
- **modelcontextprotocol/python-sdk** · clarify the validation error raised when a tool result is missing. #3053 (<https://github.com/modelcontextprotocol/python-sdk/pull/3053>)
- **confident-ai/deepeval** · format the TypeScript sources to clear a red main branch. #2823 (<https://github.com/confident-ai/deepeval/pull/2823>)
- **zilliztech/claude-context** · privacy-preserving search receipts. #406 (<https://github.com/zilliztech/claude-context/pull/406>)
- **sundeep-dayalan/sentient-trader** · a config-gated deterministic no-key replay mode. #53 (<https://github.com/sundeep-dayalan/sentient-trader/pull/53>)
- **dgenio/agentfence** · a deterministic large-trace baseline for the engine. #253 (<https://github.com/dgenio/agentfence/pull/253>)
- **vectorvoyager358/voxwire** · cover offline websocket recording replay. #44 (<https://github.com/vectorvoyager358/voxwire/pull/44>)
- **huangyikai05/Tabulint** · a deterministic large-workbook benchmark. #20 (<https://github.com/huangyikai05/Tabulint/pull/20>)
- **hsaghir/looplet** · a recipe for migrating a private tool loop to one outcome-grounded regression contract. #109 (<https://github.com/hsaghir/looplet/pull/109>)
- **smdysk/cc-parachute** · document the auto-compact filler strategy and cost ceiling for the compaction benchmark. #12 (<https://github.com/smdysk/cc-parachute/pull/12>)

Closed without merge

11 PULL REQUESTS · 10 REPOSITORIES

Declined, superseded, or fixed another way. They stay on the record: the count above only means something if the misses are here too.

- **pypa/setuptools** · repair a broken cross-reference that was failing the docs build. #5245 (<https://github.com/pypa/setuptools/pull/5245>)
- **PyCQA/isort** · wrap sorted literals under the black profile. #2573 (<https://github.com/PyCQA/isort/pull/2573>)
- **agronholm/typeguard** · handle annotated walrus assignment. #558 (<https://github.com/agronholm/typeguard/pull/558>)
- **langfuse/langfuse** · map GenAI input and output on AI SDK spans. #14636 (<https://github.com/langfuse/langfuse/pull/14636>)
- **modelcontextprotocol/registry** · security-scan receipt metadata. #1405 (<https://github.com/modelcontextprotocol/registry/pull/1405>)
- **github-linguist/linguist** · language support for QuantaLang. #7913 (<https://github.com/github-linguist/linguist/pull/7913>)
- **python-poetry/poetry** · generate launchers for gui-scripts in editable installs. #10968 (<https://github.com/python-poetry/poetry/pull/10968>)
- **pydantic/pydantic-ai** · add a turn index to agent chat spans. #6135 (<https://github.com/pydantic/pydantic-ai/pull/6135>)
- **modelcontextprotocol/inspector** · wake task polling on terminal notifications. #1728 (<https://github.com/modelcontextprotocol/inspector/pull/1728>)
- **modelcontextprotocol/inspector** · a cancel control for task polling. #1729 (<https://github.com/modelcontextprotocol/inspector/pull/1729>)
- **frostney/lwpt** · deterministic HTTP fetch-failure tests through a loopback archive-fetch override. #122 (<https://github.com/frostney/lwpt/pull/122>)

Link-list submissions, held apart

14 PULL REQUESTS · 2 MERGED

These add one line naming a Project Telos tool to a public “awesome” index: **2** merged, **10** open, **2** closed without merge. They are self-submissions to a list, they change no software, and they are excluded from every count above. A merged one means a list accepted a row and nothing more.

Creative work, shipped

Evidence of shipping, iterating on, and sustaining a widely used creative product over real time, not just a one-off release.

Elder ENB

NEXUSMODS · 936,657 DOWNLOADS

A graphics and lighting preset (an “ENB” preset) for Skyrim Special Edition, built and maintained for about two years, 2024 to 2026, across roughly 280 releases. It evolved through named editions, ENB of the Elders, Cyan Dreams, Moody Blues, Universal, and finally Elder ENB, the same discipline applied to this project: ship, take feedback, iterate, ship again. **936,657 total downloads, 160,939 unique downloaders, and 2,394 endorsements**, read from the public mod page on 2026-08-14 at version final-elder-9. A reading, not a running total: all three only go up, so the live page is the authority. View on NexusMods (<https://www.nexusmods.com/skyrimspedition/mods/117327>).

Graphics engines and reverse engineering

Native, from-scratch graphics work and the reverse-engineering practice it grew out of. Elder ENB seeded a line of D3D11/HLSL engines, and that game-systems work is where the binary-analysis and instrumentation skills came from.

Native graphics engine lineage

BRIDGE · JOJO · PLAYGROUND · RAW

A progression of from-scratch D3D11/HLSL frameworks that take ownership of a game's render pipeline through proxy-DLL interception and mid-frame compute dispatch: tens of effects and dozens of HLSL shaders (ACES and AgX, TAA, SSR, SSGI, GTAO, SDSM, volumetrics), an ImGui interface, CMake and vcpkg packaging, and a companion engine that bridges live game state to the GPU over read-only shared-memory IPC. This line was the primary vehicle for hands-on reverse engineering: binary analysis, runtime instrumentation, call-flow and memory mapping, and live game-state extraction.

Game-systems reverse engineering and automation

WORLD OF WARCRAFT · CS2 · PRIVATE

Reverse-engineered live game state and built real-time scripting and automation frameworks across several titles. Aurora is a multi-language combat-scripting framework (Lua, a custom DSL, and TypeScript-to-Lua) with dozens of SDK modules and production rotations. Its defensive counterpart, Warden, is a C++23 integrity and capability-assessment engine used as a build-and-break substrate for runtime integrity and anti-cheat research. Described by capability; the repositories are private.

Technical writing

Clear documentation for technical readers: setup guides, script references, and admin runbooks.

PowerShell script documentation

GET-STALEADACCOUNTS.PS1

A full reference for an Active Directory account-hygiene script: synopsis, parameters, usage examples, the annotated script, output format, and operational notes on replication lag and service accounts. Demonstrates writing that explains the why behind a tool, not just the syntax.

SaaS onboarding guide

WEBHOOK PLATFORM

A step-by-step getting-started guide for a webhook management product, walking a new user from first endpoint through transformations to going live, with a troubleshooting section. Demonstrates task-first documentation written against a developer workflow.

UX research and case studies

Research-led product writing: how real users move through a system and where it helps or gets in the way.

Transit fare-card UX case study

ORCA

A UX case study of a regional public-transit fare-card system, examining the rider experience and design decisions. Demonstrates structured UX research writing grounded in a real, public-facing service. Source is a formatted document; the public-safe subject is regional transit, no client or personal data.

Systems analysis

Operational and gap analysis of public infrastructure and large events: where a system is strained, and what would close the gap.

Public transit systems analysis

KC METRO

An analysis of a metropolitan public-transit system covering ridership, fleet, and capital-program trends from public data. Demonstrates turning open operational data into a clear read of system health.

Large-event transit gap analysis

FIFA WORLD CUP

A gap analysis of transit capacity against the demand profile of a major international event, with an operational playbook companion. Demonstrates scenario-driven capacity planning on public-interest infrastructure.

Data visualization

Turning public operational data into a readable picture.

Public transit operations dashboard

KC METRO / SOUND TRANSIT

An interactive dashboard component built on public ridership, fleet-availability, and capital-program data, with ridership trends, fleet status, and budget-variance views. Demonstrates clean, accessible data visualization driven entirely by open data.

What is not here

Several items from the working portfolio were deliberately excluded to keep this index public-safe. This list is part of the record, not a gap to apologize for.

- **Vendor-branded product documentation** (endpoint-security platform concept, task, and reference topics; sensor deployment and API guides; a documentation-restructuring proposal). Excluded because they are written as a specific security vendor's product documentation. Public-safe rewrites could be produced on request, with the vendor branding removed.
- **Compliance documentation** (access-control system security plan, incident-response and acceptable-use policies, ISO 27001 statement of applicability, SOC 2 trust-services documentation). Excluded because compliance documents carry client-specific control language and framework specifics that should not be published.
- **Federal grant material** (a federal grant application document). Excluded because grant templates and applications carry program- and applicant-specific detail.
- **A managed-security RFP response**. Excluded because it names a specific responding company, lists named key personnel, and includes business-registration identifiers. A fully anonymized version could be prepared if a proposal-writing sample is needed.
- **A security-log regular-expression document and a webhook API-reference document**. Excluded as a precaution because their contents could not be verified as free of environment-specific or client-specific detail from their stored format. They can be reviewed and added later if confirmed clean.

If you need a sample in one of the excluded categories, a sanitized, branding-free version can be produced on request. When in doubt, it was left out.

Download this page: Portfolio PDF. Printed from this page, so the two cannot drift apart.

Updated 2026-08-13. Every sample is owner-held and public-safe, and the samples carry the evidence rather than a description of it. Upstream counts are read from the GitHub API by tools/pr_census.py and gated by a test.