

# The Evidence

A System of Work field record, refreshed June 2026

Companion to "The System of Work."

This document is not a platform health report. It is the public-safe evidence layer for a personal thesis: when the work substrate is explicit, humans and AI agents can run parallel lanes, preserve intent, verify outputs, and leave behind enough provenance that the results can be inspected instead of narrated.

Field	Value
Evidence window	Oct 2025-Jun 2026 construction record
Current repo checkout	7b0c11217cc56eb9b38bfecd9b4e9a1c09479676
Full validation packet	b283f66e09c153e0cee95591ccde60e48b01a31a, generated 2026-06-10T23:49:02Z
Public bundle	/assets/system-of-work-evidence/2026-06-10/
Privacy posture	No raw private thread content; aggregate evidence kept as provenance

## 1. Why this evidence exists

The thesis names a coordination substrate: shared truth, authority, memory, intent, provenance, and governance. The evidence record asks whether that substrate is visible in the work itself.

The useful signal is not a generic productivity claim. The useful signal is the operating pattern across the whole construction: many AI collaboration threads, peak parallel commit bursts, bounded autonomous windows, source-owned governance, production-shaped output, and validation strong enough to keep the work honest.

The numbers below are observations from one construction. They should not be treated as universal benchmarks or adoption claims.

## 2. Highest-signal operating results

Signal	Value	Basis
Humans plus AI execution tools	2 + 4	Two engineers coordinating across Codex, Claude Code, Gemini CLI, and Replit
Merged PRs	1,717	Current GitHub merged PR count for Eigenly-Inc/fireweed-platform
All-ref commits	19,609	Current Fireweed all-ref commit count
HEAD commits	1,977	Current Fireweed main commit count
Active all-ref commit days	164	Unique commit dates across current Fireweed all-ref history
Peak commits in one minute	25	Lifetime orchestration peak
Peak commits in one day	1,018	Lifetime orchestration peak
Concurrent Codex sessions	14	Distinct sessions emitting events in the same UTC minute
P95 autonomous window	179.15m	Long-tail bounded-autonomy signal
Longest autonomous run	15.05h	Longest single coalesced-input Codex subagent window
Runtime services	21	Current production service directories
Bazel test targets	407	June validation packet test-target inventory
Preflight	44/44	Hard preflight checks passed with 0 failing

These are the page-level proof points. They show a work system that could coordinate simultaneous AI-assisted lanes into production-shaped output rather than a single assistant conversation or a monthly report slice.

## 3. Bounded autonomy

Signal	Value	What it shows
Single-input subagents	64.20%	Many spawned sessions ran from one coalesced input
P90 autonomous window	83.29m	Meaningful autonomous duration before needing repair
P95 autonomous window	179.15m	Long-tail autonomous duration under bounded context
Longest autonomous run	15.05h	Longest single coalesced-input subagent window in the local scan
Paired-turn runtime	115.27d	Aggregate paired-turn runtime across root and child Codex threads
Codex tool-call records	1.96M	Local tool-use record inside the lifetime telemetry corpus

The autonomy claim is bounded. It does not say the AI owned the work. It says the System of Work made it possible for humans to hold intent, architecture, review, and sequencing while agents executed constrained slices against shared source truth. The longest-run figure is reported as an aggregate duration only; no raw thread text is published.

## 4. What the system produced

Production surface	Count	Basis
Runtime services	21	Current Fireweed service directories
Apps	5	Current Fireweed app directories
Packages	18	Current shared package directories
Helm charts	25	Current tracked Chart.yaml files
SQL migrations	103	Current core infrastructure migrations
Proto files	38	Current tracked proto files
Source files counted	12,507	June System Pulse validation packet
Counted source lines	2,802,299	June System Pulse validation packet

The production shape matters because coordination only gets tested when work crosses boundaries: product intent into services, services into contracts, contracts into tests, tests into CI, CI into deployment, and deployment back into operating signals.

## 5. Verification surface

Verification surface	Value	Basis
Bazel test targets	407	June System Pulse validation packet
Rust test attributes	2,568	June static test-pattern scan
JS test calls	1,104	June static test-pattern scan
BUILD test calls	249	June static test-pattern scan
Workflow files	29	Current GitHub Actions workflow files in source
Preflight	44/44	<code>bazelisk test //:preflight</code> passed
Build fanout anchors	18	Affected-test fanout anchors measured
Build regressions	0	No affected-test fanout regressions

## 6. The System of Work substrate

The strongest June validation signal is that the operating layer is explicit enough to inventory and validate.

Artifact family	Count	Meaning
Domains	17	Areas with context, policy, and ownership
Workflows	126	Procedure specs defining governed work
Skills	83	Agent-facing surfaces projected from procedures
Agents	17	Formal domain agents with responsibility boundaries
Tools	112	Tool registry entries with adapters and side-effect class
Signals	352	Captured operating signals
Product-intent markdown files	349	Direction-to-specification corpus

This is the thesis in artifact form. The substrate is not only documentation. It compiles into skills, agents, tools, CLIs, validation reports, source-owned procedures, and repeatable handoff surfaces.

## 7. Validation that kept it honest

Validation surface	Result
Metrics refresh	PASS
Public pulse validation	OK
SoW validation	4,651 checks, 0 errors, 11 warnings
Work validation	0 fails, 95 warnings
Preflight	44 passing, 0 failing
TS Node type deps	PASS
GitHub workflow YAML	PASS

The warning counts are included for measurement honesty, not as the headline. The public page is oriented around what the System of Work produced and how the work moved; the underlying bundle keeps validation output available for readers who want the audit trail.

## 8. Public/private boundary

The public page and this PDF use public-safe files copied into the site repository:

- `public/assets/system-of-work-evidence/2026-06-10/README.md`
- `public/assets/system-of-work-evidence/2026-06-10/pulse/platform-health.md`
- `public/assets/system-of-work-evidence/2026-06-10/pulse/system-pulse.snapshot.v0.json`
- `public/assets/system-of-work-evidence/2026-06-10/observatory/latest.snapshot.v0.json`
- `public/assets/system-of-work-evidence/2026-06-10/observatory/index.v0.json`
- `public/assets/system-of-work-evidence/2026-06-10/validation/*.txt`

Private inputs were not copied into the public artifact set. Raw private telemetry, raw transcripts, private metrics snapshots, and logs remain out of `public/`. Local AI aggregate counts are retained in the bundle metadata as provenance, not as the headline proof.

Metric definitions are stated by basis. Current lightweight counters, lifetime telemetry, and the June 10 validation packet are not interchangeable measurement methods.

## 9. Interpretation

The evidence supports a specific thesis-bearing pattern:

- The work substrate was explicit enough to inventory.
- Human intent and review stayed separate from agent execution.
- Parallel AI lanes produced visible throughput without losing provenance.
- Governance was source-owned, not trapped in private memory.
- Verification made the output inspectable.
- Privacy boundaries were part of the system, not an afterthought.

That is the work being shown: a fully implemented System of Work with enough operating truth for serious readers to inspect the result.

## 10. Source index

Source	Public path
Audit README	/assets/system-of-work-evidence/2026-06-10/README.md
Public pulse report	/assets/system-of-work-evidence/2026-06-10/pulse/platform-health.md
Public pulse snapshot	/assets/system-of-work-evidence/2026-06-10/pulse/system-pulse.snapshot.v0.json
Public observatory snapshot	/assets/system-of-work-evidence/2026-06-10/observatory/latest.snapshot.v0.json
Public observatory index	/assets/system-of-work-evidence/2026-06-10/observatory/index.v0.json
Preflight output	/assets/system-of-work-evidence/2026-06-10/validation/preflight.txt
SoW validation output	/assets/system-of-work-evidence/2026-06-10/validation/sow-validate.txt
Work validation output	/assets/system-of-work-evidence/2026-06-10/validation/work-validate-ratchet-v1.txt
TS Node type dependency guardrail	/assets/system-of-work-evidence/2026-06-10/validation/check-ts-node-type-deps.txt
GitHub workflow YAML guardrail	/assets/system-of-work-evidence/2026-06-10/validation/check-github-workflow-yaml.txt