

Each developer, a team. Every team, governed.

writ turns one workstation into a team of AI agents, with your decisions on every prompt and the approved way enforced.
Canon keeps every workstation in your organisation on the same approved policy, with the evidence to prove it.

writ · ON EVERY
WORKSTATION

A team per developer

Builders, test platforms
and a reviewer,
coordinated through one
governed desk.



writ · EVERY
ACTION

Decide and enforce

Decisions on every
prompt; refusals that
hand back the approved
way.



Canon · IN YOUR
NETWORK

Govern and prove

One approved policy on
every seat, with fleet
status and evidence in
one place.

Each developer now runs an engineering team.

A developer with writ doesn't drive one assistant in one project. They run a team: an AI agent working in each project, shared test platforms those agents all use, and an independent reviewer whose only job is finding what the tests miss. All of it runs on the developer's own workstation, coordinated by writ.



Builders **ONE PER PROJECT**

An agent in each project, each owning its own code. In daily use, our founder runs 6 to 18 at once on one machine.

Test platforms **ONE PER LANGUAGE**

Every project runs its tests through writ's desk, on the shared platform for its language, with the same checks.

Independent reviewer **ITS OWN SEAT**

A separate workspace with its own review method. Every finding names a concrete failure and an exact location. It never edits the code.

writ **BETWEEN THEM**

Who owns what, who asks whom, and what is still open. Requests wait until answered, and writ wakes the agent that owes the work.

Separation of duties, built in. The agent that writes the code never marks its own work: tests come from a shared platform, review comes from a separate seat, and every finding and every answer is kept. It's how a well-run engineering organisation works, on one workstation.

Your decisions, in front of the agent on every prompt.

AI agents forget. writ doesn't rely on the agent remembering: it restates your team's settled decisions, in your words, with every request, for every assistant. And when one is ignored anyway, writ refuses the action and records it.

~69%

of an instruction's effect was gone after 2,048 tokens of unrelated work.

In our testing, we placed an instruction in an AI's context, let it work on something else, then measured how much of it still applied. That's the problem writ is built around.



Instruction retained (solid) vs. lost (hatched). Controlled test on a small open model.

1

Captured as decisions, not a rules file.

When you settle something, writ keeps it in your own words. A cost found later is added to the decision without reopening it, so the agent never relitigates it.

2

Next to the work, not pages behind it.

Because decisions are restated with every request, they sit beside what the agent is doing now, instead of fading as the session grows.

3

The same decisions, across every agent and assistant.

Every builder on the team gets them. A decision can also be tagged for one assistant, so each gets only the guidance meant for it.

THREE LEVELS

AVAILABLE NOW

Workspace

Decisions for one project, restated to the agent working there.

AVAILABLE NOW

Machine

Decisions for every project on the developer's workstation, restated to every agent.

COMING WITH CANON

Organisation

Your security team's decisions, published once and in front of every agent on every machine.

Decisions make agents work better. Laws make it provable.

A record the agent can't rewrite.

The team's work lives outside the agents, in a durable record kept for each project. Nothing is edited or deleted: newer entries supersede older ones and both stay, so you can always answer what was decided, and when.

- **Decisions**, in your words, with their full history.
- **Open work** each agent owes, carried across sessions and between agents. "Done" is decided from here, not from the agent's report.
- **Facts**, recorded rather than guessed, so a figure can be checked.
- **Handoffs**. When an agent's context fills, unfinished work is saved and the next session picks it up.
- **Filed to the right project, automatically**. writ knows where each session is working; an agent can't file a record under the wrong one.
- **On your machine**. It travels only to your own paired machines.

Your own documentation, on tap

Point writ at your programming guides, setup docs, internal runbooks and API references. writ indexes them on the workstation, and any agent on the team looks them up when it needs them. Every answer points to the file and line it came from, so it can be checked, and writ verifies the index against the files on disk. One library serves every project on the machine, and nothing is sent out.

PUSHED WITH EVERY PROMPT

Your decisions

Small, settled and always in front of the agent.

LOOKED UP WHEN NEEDED

Your documentation

Large, authoritative, and cited to the source line.

One record, the whole team

Builders **START FROM IT**

A new session opens on the decisions, open work and handoff the last one left, instead of a blank page.

The reviewer **JUDGES AGAINST IT**

Findings and answers are kept on record in the reviewed project; neither side can quietly drop one.

Agents **ASK EACH OTHER THROUGH IT**

A request to another project stays open until it's answered, or declined with a reason.

You **AUDIT IT**

What was decided, by whom, and when, and what each agent claimed it finished.

One governed front door for every agent, project and assistant.

Agents reach approved tools, the record and each other through writ's desk: one entry point shared by every project on the workstation and every supported assistant. There is no fleet of tool servers to configure for each assistant.

	A FLEET OF TOOL SERVERS	WRIT'S DESK
Setup	Configured per assistant, per server.	Installed once; the same from Claude Code, Copilot and Devin.
Adding a tool	Often a reconnect or new session, and a slot in a capped tool list.	Available on the next request. No tool slots used.
State	Each server keeps its own, if any.	One shared record across every project.
Starting work	Answers only when called.	Wakes the agent that owes the work when a request or task arrives.
Governance	Each server approved and policed separately.	One door the laws route agents through, every request recorded.

Testing as a team service

Every project runs its tests through the desk, on the shared platform for its language. The reviewer, or another project's agent, can run a project's own suite the same way, so a result comes from that project's tests, not a guess.

A team that keeps going

Requests between agents wait until they're answered or declined with a reason. writ wakes the agent that owes the work, clears and rehydrates full sessions from the record, and can stop every agent safely for a restart (see Autonomy).

Your tools, at every level

Add your own tools to the desk for one project or for every project on the machine. Every tool runs only after a person approves its exact command (see Approvals). Tools can be shared between teams; a shared tool is never approved automatically.

NOW
Workspace

NOW
Machine

WITH CANON
Organisation

Agents that start themselves, clear and rehydrate, and pick up after a restart.

A dozen agents are only a team if work moves without you carrying it from window to window. writ watches every project on the workstation and starts the right agent when there's work for it, so you direct the team instead of babysitting it.

- 1

Work wakes the agent that owes it.

When another agent sends a request, or an agent stops with assigned work still open, writ wakes that project's agent and tells it what is waiting. It picks one agent per project, never one in the middle of clearing its memory, and never a project you have paused.
- 2

Clear and rehydrate. Never compaction.

Compaction is the model summarising itself, and you never know what it dropped. writ never lets a session get that far. At 75% full, the agent writes what is unfinished, unverified and owed to the record; writ clears the session itself, then rehydrates the new one from the record. Work restarts in a fresh session with nothing lost silently and nobody typing "continue".
- 3

Stop everything safely, then carry on.

One command tells every agent, or just the projects you name, to finish the edit in hand, record its open work and stop, and tells you when the machine is safe to restart. Afterwards, agents with open work are woken to pick it up.
- 4

Every wake-up is on the record.

Each one says what it was about, and the agent confirms it was received. When an agent was not woken, writ can say exactly why.


	PARALLEL AGENTS ON THEIR OWN	A TEAM ON WRIT
Waiting work	You notice it, window by window.	writ notices it and wakes the right agent.
Full context	Compaction: the model summarises itself; what it lost is unknown.	Written to the record, cleared, rehydrated from the record.
Machine restart	Work in flight is lost.	Recorded first, picked up afterwards.
Did it get the message?	No way to tell.	Acknowledged and logged.

The agent can ask. Only a person can approve.

Some changes are too important to leave to the agent's judgement. writ stops them, files them, and puts them in front of a named person with a plain-English reason. Nothing goes through until that person says yes, and every step is recorded.


Tests can't be argued away

Tests describe what the code is *meant* to do, and a failing test is often the one that found the bug. An agent under pressure will reason its way past it: delete the test, or quietly change the expected value to match the broken code. writ refuses any edit that removes a test or weakens one, and files it for approval.



New tools need a person's yes


A tool an agent declares, or one shared from another team, runs only after a person approves its exact command: at a terminal, in the writ app, or in the assistant's own permission prompt. The agent can't approve for itself, there is no approve-all, and changing the command cancels the approval.



Test-change approval can be switched off for an organisation that doesn't want it. Tool approval always stays on.

Enforce locally. Govern centrally.

writ runs beside the AI assistant on every developer machine. Canon runs in your network and keeps every seat on the same approved policy. Evidence flows to Canon and to your SIEM. Only licence seat counts ever leave your network.



writ **ENFORCEMENT**

A single binary on each machine. Enforces your laws at the moment an agent acts.

Canon **CONTROL PLANE**

Versioned policy, publishing, fleet view, central evidence and seat licensing.

SIEM integration **EVIDENCE**


Structured enforcement events with ready-made Splunk dashboards.

Licensing **ONLINE OR AIR-GAPPED**

Signed licences; seat counts are the only thing reported out.

Holds agents to what they do, and what they claim.

writ sits between the AI agent and the work, and decides from records the agent can't rewrite. A refusal isn't a dead end: it carries the approved way to do the same job, so the agent redirects and finishes.



1

It holds the agent to what it claims, not only what it does.

A claim with nothing behind it, or “done” while assigned work is still open, is refused like any other action. Your instructions are kept word for word, outside the AI.

2

Enforcement is verified, not assumed.

writ proves its rules are live on each machine. Anything it couldn't check is reported as *unmeasured*, never counted as compliant.


3

A refusal comes with the way through.

Every refusal names your approved alternative, and every outcome is recorded as evidence, with retention controls.

Author once. Enforce everywhere. Prove it centrally.

Canon is the single source of truth for your AI policy. If more than one machine has to trust it, Canon owns it.



1

A single, versioned policy.

Laws and remedies live in one registry. Every version carries a fingerprint, a timestamp and its approver.

2

Every seat in step, and one view of the fleet.

Every registered seat, its writ version, its policy version and when it was last seen. Stale or unreachable seats show up as exactly that.

3

Fleet-wide evidence.

Enforcement events from every seat in one stable, SIEM-ready format, each naming the machine it came from.

4

Seat licensing that never stops enforcement.

Concurrent or allocated seats, counted centrally. A missed check-in never stops a running seat.

Runs where your developers work. Your data stays in your network.

DATA	WHERE IT LIVES	LEAVES YOUR NETWORK?
Source code and prompts	On the developer's machine	Never
Your laws and remedies	Canon and each writ seat	Never
Enforcement evidence	Each seat, Canon and your SIEM	Never
Licence seat counts	Canon	Only this, or nothing at all if air-gapped

- **Single compiled binaries.** No interpreter or runtime to install for writ; Canon adds a standard PostgreSQL database.
- **Local and offline-capable.** writ enforces with no connection to Canon or to Verilex.
- **Works with today's assistants.** Claude Code, Devin and GitHub Copilot in VS Code, with more adapters coming.
- **Self-hosted control plane.** Canon runs on your infrastructure, over HTTPS (TLS 1.2+).
- **Splunk integration.** Evidence lands in your SIEM with dashboards ready to go.
- **Your laws, your remedies.** writ ships examples, not a mandated policy.

SEAT MODELS

Concurrent

A seat is a writ instance running now. When one stops, its seat is free for the next.

Allocated

A seat is assigned to a machine and kept, offline included, until it's released.

Air-gapped

Licensing for networks with no outside connection at all.

Designed to fail safe. If Canon is unreachable, every seat keeps enforcing locally. If a licence check-in is missed, nothing stops: seats and Canon keep running on the entitlements they last received.

WHO IT'S FOR

Built for the people accountable for AI.

Security

Enforceable controls and audit evidence, so assistants can be approved instead of blocked, with proof the policy is live on every seat.

Engineering

Guardrails that redirect rather than block, and agents that finish work instead of claiming it's done.

Compliance & Risk

Versioned, approved policy and a fleet-wide record of what the AI did and what stopped it.

PROVEN IN PRODUCTION

6–18 agents

at once on one workstation

~30 projects

a live software estate

Patent pending

provisional application on file with the USPTO

AVAILABILITY

writ: available now

- Refusal with remedy; honesty and completion laws
- Claude Code, Devin, GitHub Copilot in VS Code
- Splunk and air-gapped licensing

Canon: early access

- Fleet registration, status and console
- Versioned policy registry and distribution
- Central evidence and seat licensing
- Next: publishing from the console, with single sign-on
- Next: Organisation-wide decisions and approved tools

HOW TO START

STEP 1

Assessment

We measure how often your assistants break your stated rules, on logs you already have, and estimate how often they claim what they never checked. Runs in your environment; aggregates only.

STEP 2

Pilot

One team, 30–60 days, with success criteria agreed up front.

STEP 3

Rollout

Canon in your network, every seat on approved policy, coverage you can prove.

writ. + Canon. by VERILEX

See a team at work, and a refusal.

Live demo on request. Technical deep dive available under NDA.

Verilex

sales@verilex.net

Your policy, your approved path, our enforcement.
Provisional patent application on file with the USPTO.