

The Novel Writing Machine

A System Documentation — v0.2

An agentic system for long-form fiction organized like a Formula 1 race weekend. Fifteen specialized roles protect author judgment at the moments where it actually matters.

ELIZABETH ANN WEST

AGENTIC-SYSTEMS

LONG-FORM-FICTION

AI-AUTHORSHIP

WORKFLOWS

SYSTEM-DESIGN

Abstract

Single-agent pipelines fail at long-form fiction in predictable ways: physical facts drift across chapters, point-of-view characters leak foreknowledge they have not yet earned, and the prose flattens into a generic, recognizably-AI voice. This paper documents a different approach — an agentic system that organizes the writing of a novel like a Formula 1 race weekend, splitting the work across fifteen specialized roles arranged into three gated phases. A **Garage** builds the foundations before a word is drafted; a **Race** writes the book scene by scene under constant audit; a **Parc Fermé** inspects the finished manuscript from altitude. The load-bearing mechanism is a per-character, per-scene knowledge state that makes foreknowledge leaks structurally impossible, paired with a five-level severity system that reserves the author's judgment for the four decisions that actually matter and lets the machine handle the rest. The result is a workflow in which an author can ratify a night's setup, walk away, and wake to a queue of scenes drafted in her own voice — with a complete log of every call made along the way.

Single-agent writing pipelines break in predictable ways.

Long-form AI generation produces three failure modes that no amount of prompt-engineering can quietly fix.

Eye colors drift. A character whose eyes were dark in chapter three are hazel by chapter twelve. The model holds a vague impression of the character, not the fact, and the fact mutates over thousands of generated tokens.

Characters know things they shouldn't yet. The model has read the whole outline and leaks foreknowledge into a POV character's thoughts. Elizabeth Bennet references Wickham's elopement in chapter ten — ten chapters before she is told. The reader catches it. The illusion breaks.

The prose flattens into a generic competent voice. Sentence rhythms converge. The same constructions recur with quiet inevitability — *"It wasn't X but Y."* *"The air shifted."* Names cluster around an unmistakable AI-default palette: Marcus Chen, Sarah Whitfield, Elena Ashford. The book becomes legible as AI-written even when it isn't trying to be.

A single model writing a 90,000-word novel is a single human writing a 90,000-word novel without sleep, notes, or judgment. The structure is what's missing, not the talent.

The machine in this document is the structure. It splits the work across many specialized roles, gives each one a narrow lane and clear constraints, and stages their cooperation as the three acts of a Formula 1 race weekend: pre-race setup, the race itself, and post-race inspection.

Three levels. Each one a gate the work must pass through.

Every Formula 1 weekend has the same architecture. The car gets built in the garage — engineers tune setup, mechanics check tolerances, the team confirms scrutineering before the car ever rolls onto the track. The race itself runs lap by lap, with stewards watching every move and a team principal calling strategy from the pit wall. After the chequered flag, the cars go into parc fermé — the impound area — where inspectors check that everything that crossed the finish line meets specification.

A novel passes through the same three phases.

PHASE	F1 ANALOGY	WHAT HAPPENS
Garage	Pre-race setup	Dossier, outline, world-state, characters, voice anchor built before a word of prose is written
Race	The race	Prose drafted scene by scene; stewards audit every lap; team principal handles flags
Parc Fermé	Post-race inspection	Six inspectors sweep the completed manuscript; flagged candidates surface to author for ratification

Two more pieces of vocabulary travel through the whole system. **Scenes are laps** — the atomic unit of story, small enough for one agent to hold complete context. **Chapters are stints** — continuous runs of laps on the same set of tires, between pit stops.

The dashed red path that most novel pipelines miss: when the prose reveals something better than the outline planned, the system routes back to the garage — not as a failure, but as a discovery. The team principal makes the call. The outline updates. The race resumes from the affected scene with the new direction locked in.

The three-level architecture

Every novel passes through three gates — and routes back to the garage when the prose beats the plan

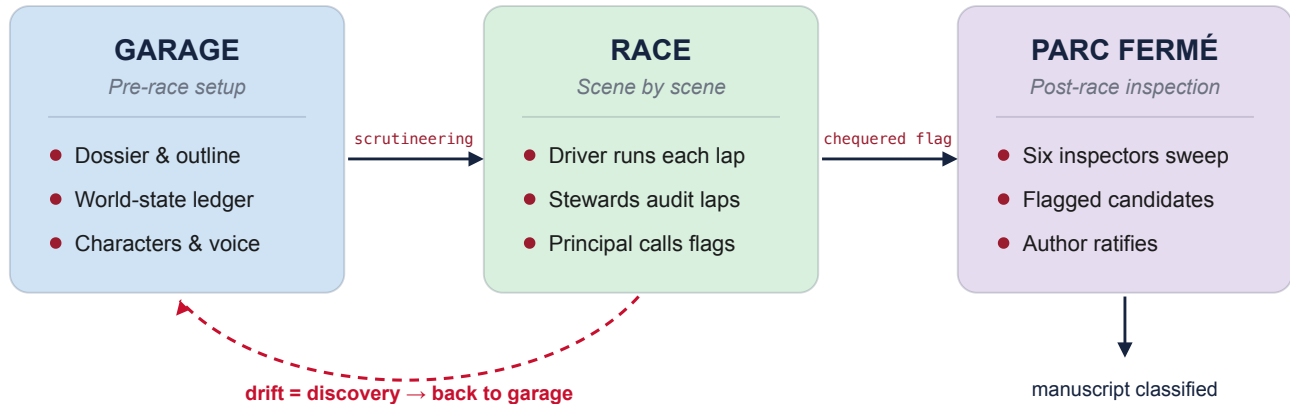


Figure: the three-phase pipeline — Garage, Race, Parc Fermé — with the discovery loop that routes a scene back to the garage when the prose outgrows the outline.

● SECTOR 03 · GARAGE

Everything before a single word of prose is written.

Five skills work in sequence to produce the foundations the rest of the system depends on. The gate at the end is called **scrutineering**: the car cannot roll onto the track until every deliverable passes inspection.

The garage exists to make the most expensive failure mode of long-form generation impossible: writing 40,000 words before noticing the structure doesn't hold. Each garage skill produces a concrete artifact, and the next skill reads the previous one's output.

SKILL	TYPE	DELIVERABLE
Story Engineer	Creative	17-section dossier, scene-by-scene outline, payoff architecture, tension curve, telemetry baselines
World Engineer	Creative	World-state ledger with static facts and per-character knowledge-state schema
Character Lead	Creative	Interiority files: backstories, micro-histories, sensory associations, body memory, fears with origins
Naming Mechanic	Creative	Cast names from period/region-appropriate pools with explicit anti-cluster filtering — no Marcus Chen, no Whitfield, no Ashford
Style Auditor	Judgment	Voice anchor built from author's prior work; manuscript-specific crutch list with local density limits

The knowledge-state machinery

Buried inside the World Engineer's output is the single most important piece of machinery in the whole system: a per-character, per-scene knowledge state. Every fact in the ledger carries a `known_to` field listing which characters know it, and a `known_since` field marking the scene where they learned it.

Before the Driver writes any scene, it queries the ledger: *what does Elizabeth Bennet know as of scene fourteen?* The query returns a scoped subset of the ledger — not the whole world model, but only the facts that character can reference at that moment in the story. The leak that produces "characters knowing things they shouldn't" stops being possible, because the Driver never sees the facts it isn't allowed to use.

The exit gate at the end of the garage is called **scrutineering**. The team principal reviews every deliverable — dossier, outline, ledger, interiority files, name list, voice anchor, crutch list — and refuses to enter the Race until each one passes. There is also an optional **qualifying lap**: one pilot scene drafted before the full race, so any problems with voice or interiority surface before twenty more scenes get built on top of a bad foundation.

Each scene is a lap. Each lap gets watched.

Drafting happens one scene at a time. The Driver runs the lap through a seven-step loop. While it drafts, the Pit Crew preps for the next. After the lap, four stewards inspect what was written.

The seven-step loop

01 Gather

Context from the dossier and world-state ledger scoped to this scene — including the per-character knowledge query.

02 Plan

Scene-level intent: what needs to happen, what the POV character knows and feels, where tension needs to land.

03 Draft

First-pass prose. Full scene, written in the author's voice anchored from the Style Auditor's file.

04 Repetition pass

Reads the lap against the previous five scenes, hunting for echoes: repeated body language, recycled metaphors, opening structures that pattern-match across chapters.

05 Continuity pass

Cross-checks the draft against the world-state ledger. Flags any reference to a fact the POV character does not yet know at this scene.

06 Craft pass

Where the prose becomes prose. Named beat identified; every line measured against it; lines that do no work killed or recast.

07 Final draft

Applies all three editorial passes and outputs the ratified scene.

The seven-step lap

How the Driver writes one scene — three editorial passes built into every lap



Figure: every lap runs the same seven steps; steps four through six are the editorial passes — repetition, continuity, and craft — that catch drift before a scene is ratified.

The Pit Crew

While the Driver runs a lap, three pit-crew mechanics work in parallel on the work that surrounds the prose. They are the difference between an autonomous race and a stop-and-start race.

MECHANIC	FUNCTION
Backstory Mechanic	Generates interiority on demand for minor characters who acquire significant page time mid-novel — the cast members the Character Lead deliberately did not pre-build
Continuity Mechanic	Writes updates to the world-state ledger after every lap — new knowledge, location changes, prop moves, relationship shifts. The only agent with write access to the dynamic ledger
Sidethought Catcher	Captures emergent ideas the Driver throws off mid-scene — things that don't belong in this lap but might be the seed of a future one. Filed in the race log

The Stewards Panel

Four jurisdictions, each scoped narrowly, each silent outside its lane:

STEWARD	JURISDICTION	CATCHES
Continuity Steward	World-state ledger	Foreknowledge leaks, physical drift, spatial impossibilities, prop errors, timeline contradictions
Voice Steward	Voice anchor (craft-editor in audit mode)	Voice divergence, crutch patterns, generic language, style drift from the author's anchor
Character-Truth Steward	Interiority files	Behavior inconsistent with character's operating system, dialogue that doesn't match private vocabulary
Pacing Steward	Tension curve and telemetry	Missed word-count targets, divergence from target tension, missed axis movement, energy-curve anomalies

The severity gradient

Each steward assigns severity to its flags. The Team Principal consolidates the panel's reports and decides the action. There are five levels:

The severity gradient

Five levels of steward flag — from a logged note to a full restructure

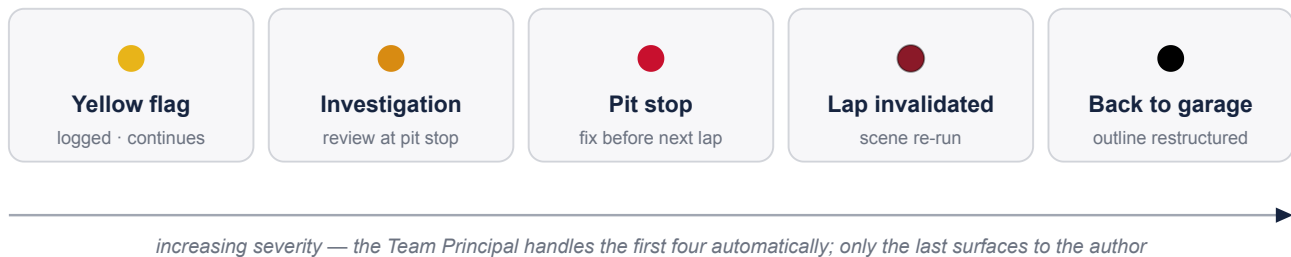


Figure: the five-level severity gradient. Yellow flags log and continue; a back-to-garage call is the only level that interrupts the author for a decision.

FLAG	WHAT HAPPENS
● Yellow flag	Note logged in the race log. The lap counts. The race continues.
● Investigation	Flagged for review at the next pit stop. The lap counts provisionally; the race continues.
● Pit stop	Come in now. The specific issue gets fixed before the next lap. The lap counts after the fix.
● Lap invalidated	The Driver re-runs the scene with the steward's note as a constraint. The old lap is discarded.
● Back to garage	Structural problem the race can't fix. Outline updates, ledger updates, then the race resumes from the affected scene.

Drift = tangent gets corrected. Drift = discovery rewrites the outline. The author makes that call.

The Team Principal handles most calls automatically. Yellow flags log and continue. Investigations queue for review. Pit stops fix and resume. Lap invalidations re-run the scene with new constraints. Only the back-to-garage class of call — and specifically the distinction between drift-as-tangent and drift-as-discovery — surfaces to the author for ratification.

One craft authority, four scopes. The *craft-editor* skill is the one place line-editing judgment lives. The Driver calls it at Step 6 of every lap. The Voice Steward calls it again in audit mode on the finalized lap. Parc Fermé calls it a third time across the whole manuscript. And it can be called standalone whenever the author wants a clean line edit on any passage. Same philosophy every time — different scopes, one source of truth for prose quality.

● SECTOR 05 · PARC FERMÉ

After the chequered flag, the inspection.

Stewards catch what's visible lap by lap. Parc fermé catches what's only visible from altitude — patterns across the whole manuscript that no single scene reveals.

Six inspectors run in sequence after the final scene completes. None of them edit the manuscript directly. Each produces flagged candidates with reframes, and the author ratifies what gets applied.

01 **Craft Editor (manuscript scope)**

The primary book-wide craft sweep — the same skill the Driver and Voice Steward use, now applied to the whole manuscript at once. Weight, voice, vague nouns, lived-vs-written. Line-level rewrites surfaced as candidates for ratification.

02 **Grammar Mechanic**

Subordinate to craft-editor. Sentence rhythm and opener variety across chapters — runs of uniform sentence lengths, repeated opening structures, mechanical tics that only show up at book scale.

03 **Crutch Inspector**

Subordinate to craft-editor, sharing the same avoid-list. Global frequency of crutch patterns. The metric is density and frequency, not presence: twice in a chapter is fine, three times in two paragraphs is not.

04 **Pacing Inspector**

Compares actual scene-by-scene tension and word counts against the target tension curve. Surfaces drag zones and energy dips.

05 **Reader Simulator**

Sequential first-read of the manuscript reporting engagement state per chunk — confused, bored, hooked, moved, suspicious. The curve matters more than any single moment.

06 **Payoff Auditor**

Cross-checks the manuscript against the payoff architecture. Every setup pays off. Every payoff has its setup. Promises made are promises kept.

Each Parc Fermé flag carries a location, an observation, a severity, and one to three reframe candidates. The author reads the full flag list and decides per flag: **accept** (apply the reframe), **reject** (mark it as kept)

intentionally), or **revise** (write the fix herself). Ratified edits apply. The manuscript is then classified — race result official.

● SECTOR 06 · THE ENTRY POINT

Every race begins with a book order.

The author pitches the book the way an editor would brief a developmental team. Hook, audience, payoffs, voice references. Concrete enough to constrain. Open enough to discover.

The system runs on a single file at the root of each book's project: `book-order.md`. The payoffs section is the load-bearing part. Generic payoffs produce generic books.

```
# Required payoffs – what the reader needs to feel by the last page

payoff_1: the heroine forgives herself for the stable boy incident
payoff_2: the hero earns the heroine's trust without an apology speech
payoff_3: the magic system's cost lands as devastating, not cool

# Voice references – books whose voice samples the Style Auditor reads
voice_reference: [prior_book_1, prior_book_2]

# Disclosure
disclosure: "This book was written with author-directed AI tools."
```

The book order goes into the Team Principal's hands. The Team Principal reads it, refuses to start if any required field is still bracketed, prints the planned phase sequence, and waits for the author to confirm before kicking off. From there, the machine runs.

Autonomous on the boring parts. Gated on the parts that matter.

The machine is not designed to produce a finished manuscript without the author. It is designed to remove the work that doesn't need authorial judgment, so the work that does can get full attention.

Four pause points exist by design. They are the places where, if the machine pushed through alone, the manuscript would land in revision territory anyway.

THE AUTHOR RATIFIES	THE MACHINE HANDLES
Scrutineering — Garage outputs feel right before the race begins	Yellow flags — Logged and continued automatically
Naming lock-in — Names are identity; confirm before they freeze	Pit stops — Self-resolved before the next lap
Drift = discovery — The moment the book becomes itself	Lap invalidations — Re-run with new constraints
Parc fermé flags — Every reframe is a candidate, not a command	Continuity bookkeeping — The world-state ledger updates after every scene

What this enables: ratify the garage outputs in an evening session. Launch the race. Walk away. The machine drafts overnight — scene after scene, stewards flagging, the team principal handling most calls automatically. In the morning, the author wakes to a queue of new scenes drafted in her voice, a race log of every call along the way, and a short list of decisions waiting on her judgment.

For an 85,000-word novel, that is a week of those sessions rather than one overnight. But "wake up to eight new scenes drafted" stops being a dream and becomes a workflow.

Every skill. Every status. One table.

A complete enumeration of the racing team, what each role does, and what's built versus still in design.

Status legend: BUILT = complete SKILL.md, ready to run. WRAPS = calls an existing foundational skill.

STUB = role scaffolded, not yet fully designed.

Sector 01 — Garage

SKILL	STATUS	DESCRIPTION
Story Engineer	WRAPS	Dossier, scene-outline, payoffs, tension curve, telemetry. Wraps story-dossier-cascade.
World Engineer	BUILT	World-state ledger plus knowledge-state machinery. Static facts and dynamic state schema.
Character Lead	BUILT	Interiority files for every major character. Off-screen gravity wells the Driver writes from.
Naming Mechanic	STUB	Cast names from period/region-appropriate pools, anti-cluster filtered.
Style Auditor	STUB	Reads voice samples, builds the voice anchor and the manuscript-specific crutch list.

Sector 02 — Race

SKILL	STATUS	DESCRIPTION
Team Principal	BUILT	Orchestrates the whole pipeline. Makes every severity call. Only agent that restructures the outline.
Driver	BUILT	Runs one lap through the 7-step loop in scene mode. Wraps chapter-writer-7step.
Backstory Mechanic	STUB	Generates interiority on demand for minor characters who acquire significant page time mid-novel.
Continuity Mechanic	STUB	Writes updates to the world-state ledger after every lap. Only agent with write access.
Sidethought Catcher	STUB	Captures emergent ideas the Driver throws off mid-scene. Filed in the race log.
Continuity Steward	STUB	Reads against the world-state ledger. Flags foreknowledge leaks and drift.
Voice Steward	STUB	Runs craft-editor in audit mode against the voice anchor.
Character-Truth Steward	STUB	Reads against the interiority files. Flags behavior inconsistent with character's OS.
Pacing Steward	STUB	Reads against the tension curve. Flags missed targets and energy anomalies.

Sector 03 — Parc Fermé

SKILL	STATUS	DESCRIPTION
Craft Editor	BUILT	Primary craft sweep at manuscript scope. Called from four contexts: Driver Step 6, Voice Steward, Parc fermé, standalone.
Grammar Mechanic	BUILT	Subordinate to craft-editor. Rhythm runs and opener variety across chapters.
Crutch Inspector	BUILT	Subordinate to craft-editor. Crutch frequency and density, book-wide.
Pacing Inspector	STUB	Compares actual tension and word counts against the targets.
Reader Simulator	STUB	Sequential first-read engagement curve.
Payoff Auditor	STUB	Setup ↔ payoff cross-check.

Foundational — Wrapped

SKILL	STATUS	DESCRIPTION
chapter-writer-7step	EXISTING	The 7-step loop the Driver wraps. Standalone for chapter mode; scene mode for the Race.
craft-editor	BUILT	The one craft authority. Called from four scopes.
story-dossier-cascade	EXISTING	The 17-section EAW worksheet. Includes the Narrative Physics Engine — tension axes, thresholds, entropy spikes, the 11 Laws.
chapter-word-counter	EXISTING	±200 word tolerance checks. Used by the Driver and the Pacing Inspector.

Architect: Elizabeth Ann West · Design partner: Claude (Anthropic) · Filed: May 2026 · v0.2

Status: Race phase implemented, Parc fermé stubs next. 7 skills fully built, 8 stubs to design.