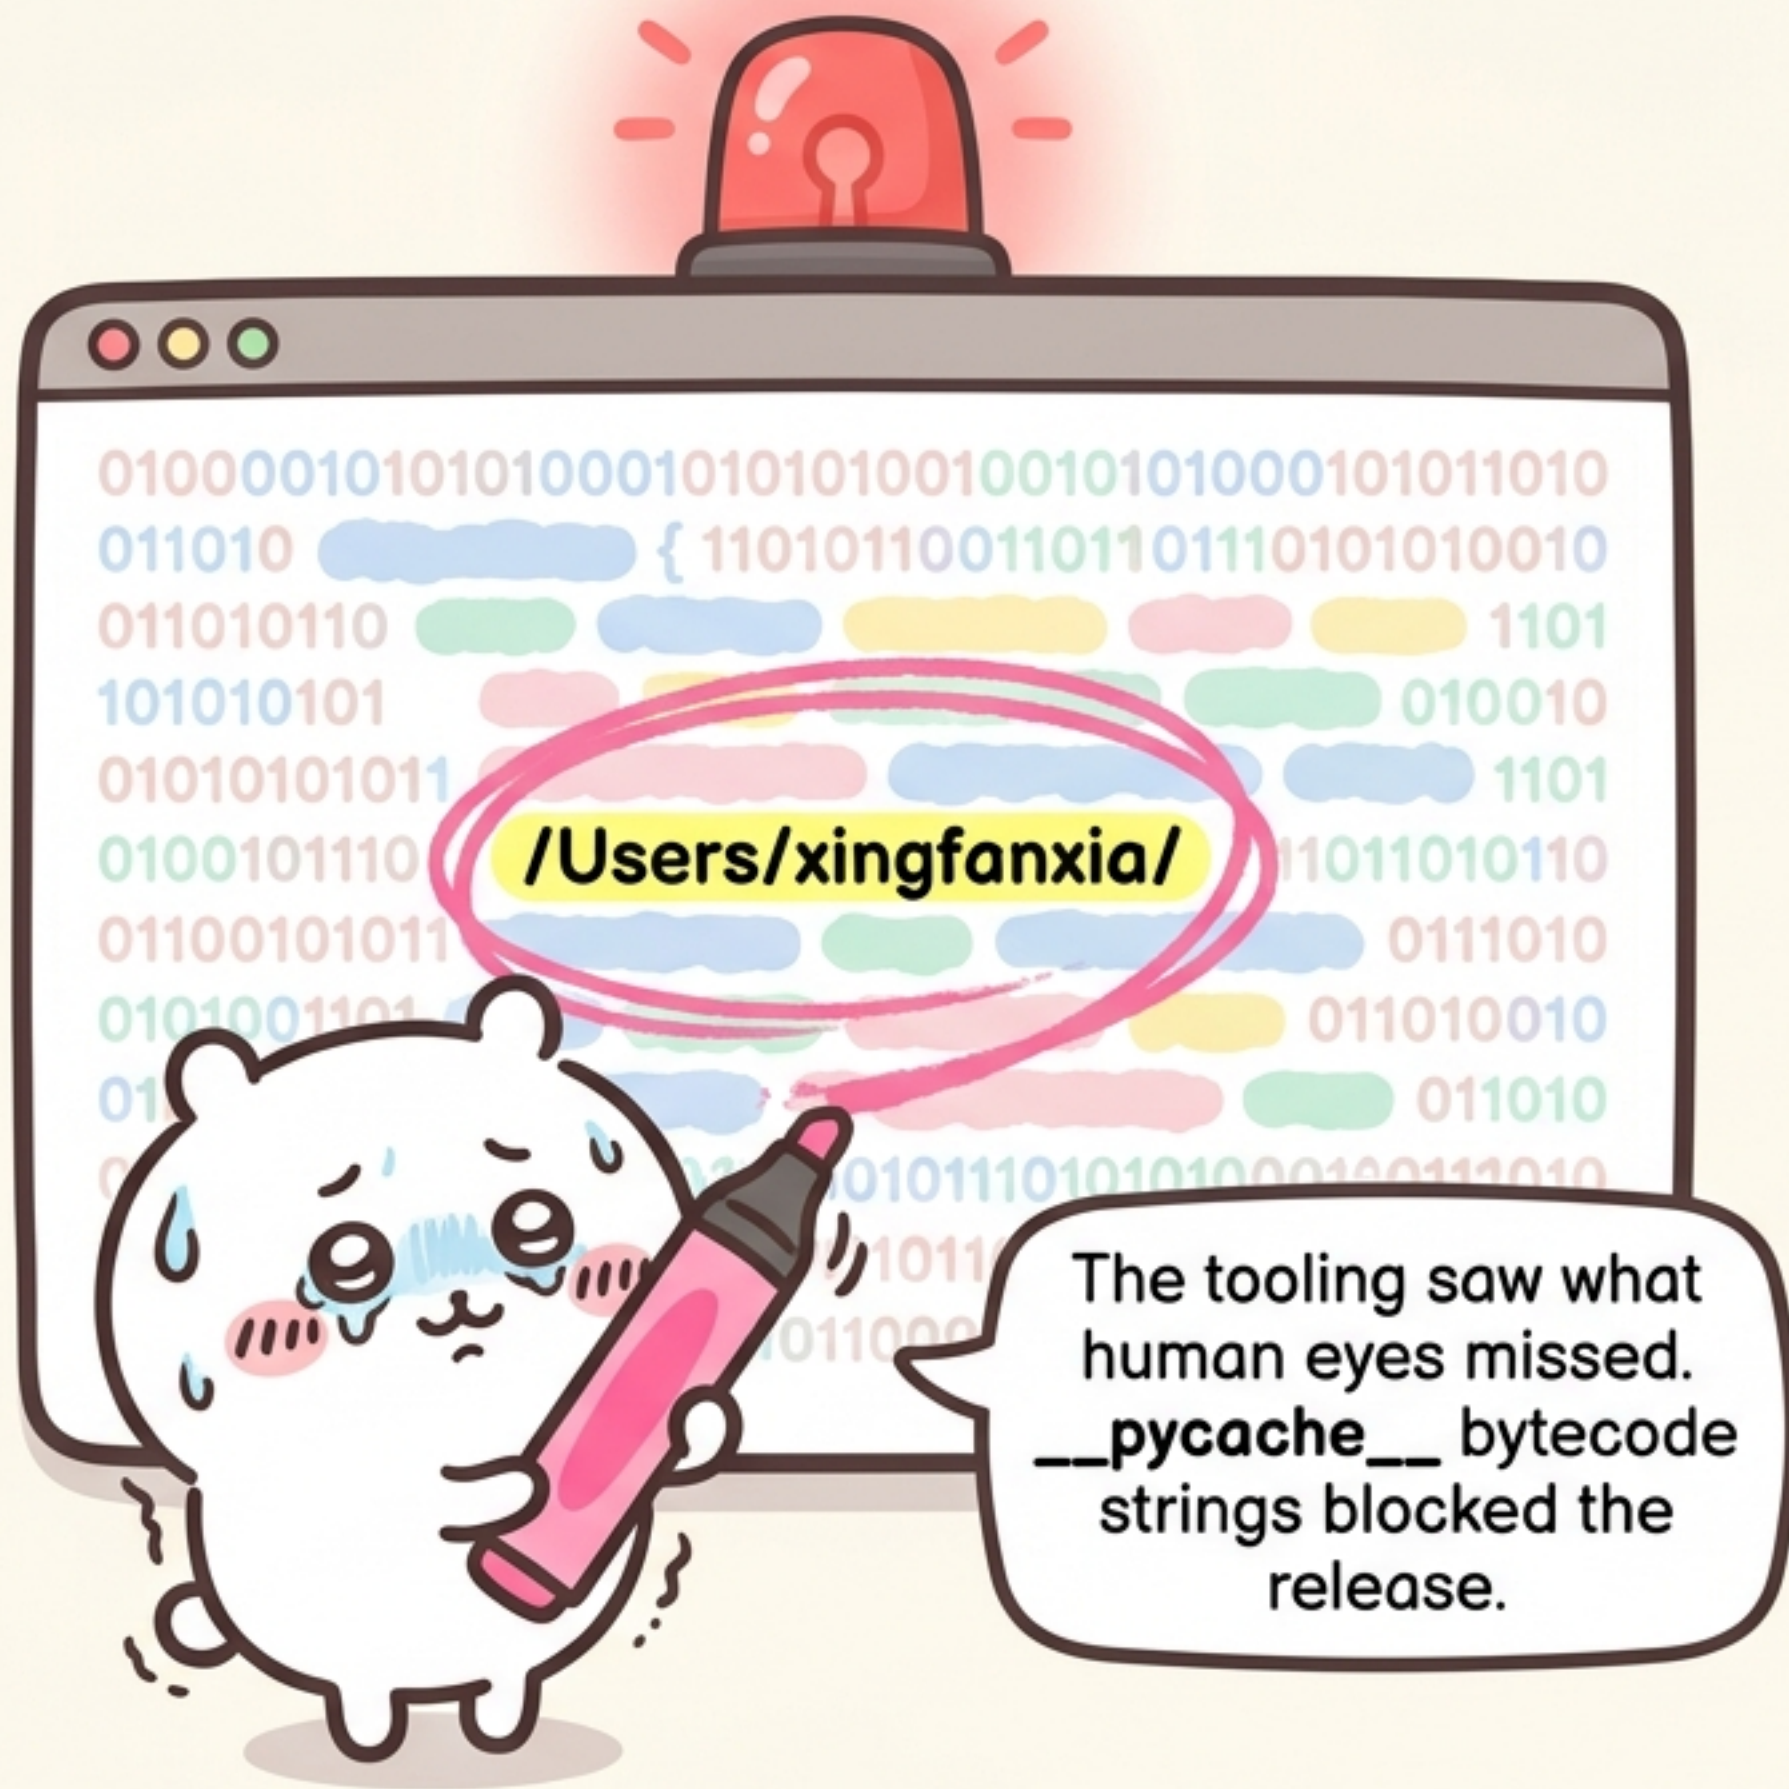





The Trigger: A Fail-Closed Safety Net



Build the safety net. Trust the safety net.
A custom configuration is interesting, but an aggressive audit-and-refuse mechanism is essential to prevent self-inflicted leaks.



Not Just Code: Categorizing 50+ Custom Skills

The Tool Typology

The Engine (Engineering)

- **audit-fix-loop:** Iterative review until code is clean.
- **pr-fix-loop:** Autonomous response to PR feedback until CI is green.
- **codebase-sweep:** Full repository audit.

The Polish (Design & Writing)

- **khazix-writer:** 4-layer self-checking Chinese narrative prose.
- **huashu-design:** Hi-fi prototyping with 20 baked-in design philosophies.

The Soul (Weird & Personal)

- **banxian-skill:** 3-system divination. Python generates hexagrams; AI interprets.
- **jewelry-marketing:** Xiaohongshu copy generated from a single product photo.

The weird skills are the soul of the config.
It is a highly opinionated, personal digital workshop—not a generic developer toolkit.

Orchestration: Parallelism & Risk Routing

pet-forge: Custom desk-pet animations running alongside serious workflows.



Risk Router
(big-task skill)
Routes by signal
(auth, schema), not
by file count.

Discuss → Plan → Execute

Forces architectural
decisions before
execution starts.

gsd-plan-checker
(Goal-backward verification)

**code-reviewer &
security-reviewer**
(Dedicated parallel
analysis)

neat-freak
(Enforces docs hygiene;
stops drift between
code and CLAUDE.md)

1M context windows hold the noise. Don't spawn subagents to "protect" context; spawn them for genuine parallel execution. One sentence triggers a massive, autonomous workflow.

Quality Control: Banishing AI Slop



Never trust a single-pass self-review.
Force agents to critique their own plans through three distinct lenses. An AI's first draft will self-assure; its third pass will find the P0 bugs.

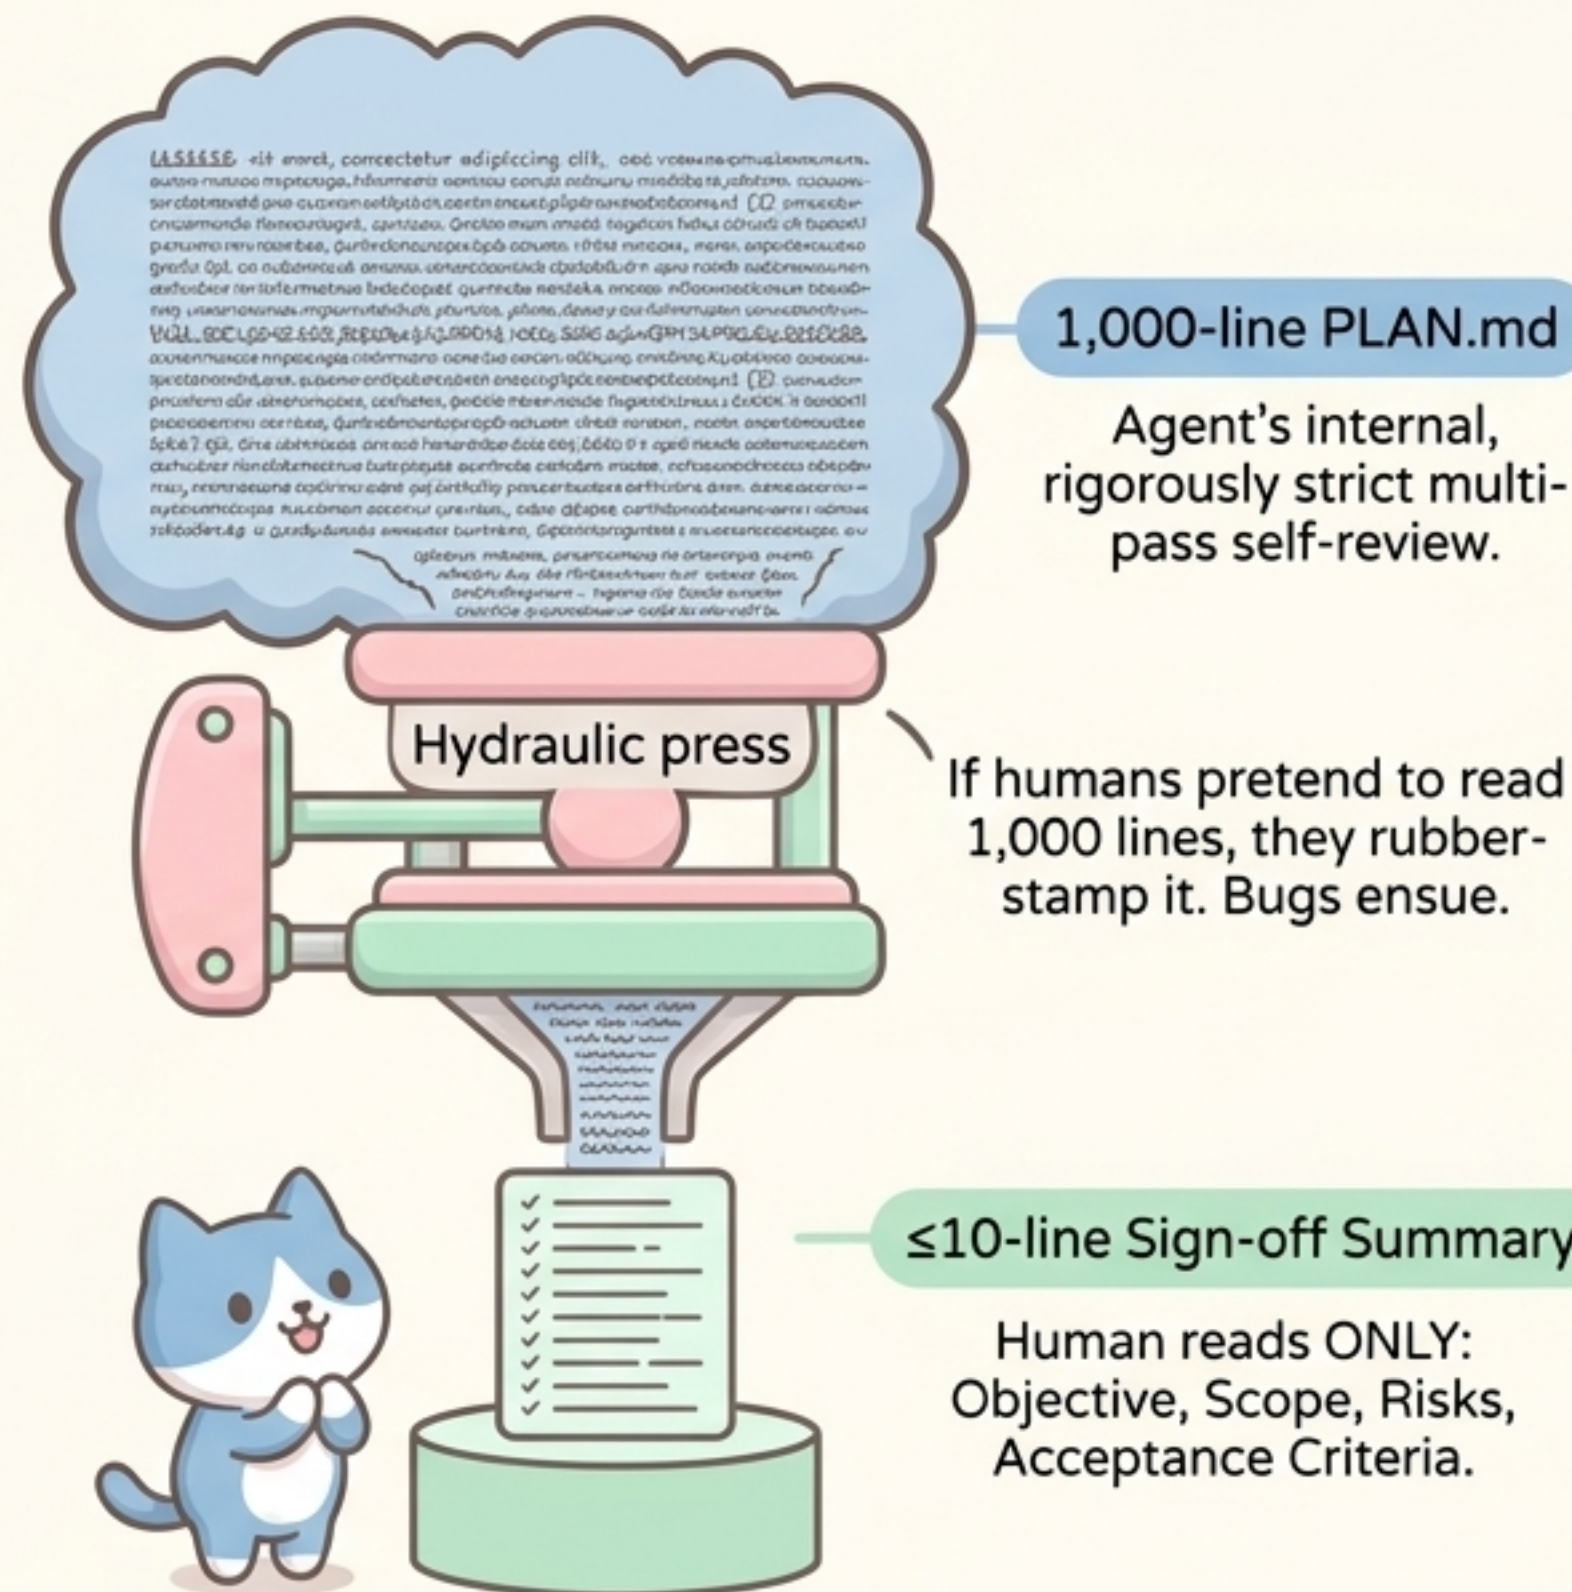
The Anti-Slop Filter

- ☒ **Pass 1: Architecture**
Does the design actually work?
Demand empirical grep verification before trusting claims about file content.
- ☒ **Pass 2: Security**
Are failure modes enumerated?
Are industry-standard checks (like gitleaks) present?
- ☒ **Pass 3: Human Factors**
Will a real human actually use this without footguns?

Extreme Ownership: “AX Doesn’t Review”



Once you admit
you won't read a
massive plan, you
force the system to
filter it for you.



The agent knows there is no human safety net, so its own self-review gets stricter. Your job isn't to read the code; your job is to review the executive summary and approve the launch.