


Agentic AI Economics: from token burn to business value.

AI agents are already committing, reviewing and merging code — Copilot, Claude, Cursor, Code-Rabbit, Qodo, Lovable and more. Most organizations can answer neither the risk question (**which agents can change our software, with what privileges?**) nor the economics question (**what does the spend and token burn actually return in accepted code and delivered features?**). BlueFlag's AI Insights answer both from one place: every agent discovered and governed as an identity, and every dollar of AI spend traced to accepted code.



All figures from a live example environment (90-day window) — shown to illustrate the product; customer values are produced by a read-only scan.

EXECUTIVE SUMMARY

Agents are identities. Spend is a control. Both need governance.

The AI Insights module treats every coding agent as a first-class identity — discovered, attributed, access-reviewed and scored — and treats AI spend as an engineering-economics dataset: seats vs. active users, spend vs. tokens vs. lines accepted, model adoption across the estate. The result is one picture that serves the CISO (risk), the CFO (cost) and the CTO (velocity).

WHAT THE MODULE ANSWERS

- **Which AI agents exist** in our SDLC — active and shadow — and in which repositories?
- **Who wrote what:** agent-primary vs. human-primary vs. co-authored commits and PR activity
- **What does AI coding cost** — and what comes back as accepted code?
- **Which models and frameworks** are actually in use, cloud-managed or self-hosted?

THE QUESTIONS LEADERSHIP IS ASKING

- › Can an unregistered AI agent merge to production — today?
- › Why do we pay for seats nobody uses while unlicensed tools burn tokens?
- › What did that feature actually cost in AI spend?
- › Where does AI-generated code enter our regulated applications?

THE FEATURE MAP

Every AI Insights feature, mapped to risk removed and value created

FEATURE	WHAT YOU SEE IN THE PRODUCT	RISK REMOVED · BUSINESS BENEFIT
SDLC Agent Activity	Active agents ranked by real activity — commits, PR comments, reviews (e.g. Code-Rabbit 4,358 events, Qodo 238, Cursor 76) with 90-day trend charts.	Risk: ungoverned agent write activity. Benefit: know which agents actually work in the factory — and which are noise.
Shadow-AI / Agent Discovery	Discovered agents the organization never registered — with number of identities and repositories each touches (e.g. Claude on 4 identities / 4 repos, Lovable 1 / 0).	Risk: Shadow AI — unmanaged agents and tokens with code access. Benefit: complete inventory before governance; nothing changes code unseen.
Agent vs. Human Attribution	Every commit split into agent-primary, human-primary and co-authored; PR opened / approved / commented / merged per agent.	Risk: unclear provenance of AI-generated code. Benefit: audit-ready authorship trail; review effort focused where agents write.
License Utilization	Active users vs. licensed seats per tool over time — including usage with zero seats (e.g. Claude Code: 4 active users, 0 seats).	Risk: unlicensed / shadow usage outside procurement. Benefit: right-size seats, recover unused licenses, formalize what teams already use.
Usage Activity & AI Spend	Total AI spend with per-day trend, switchable between Spend, Lines accepted and Tokens — per tool and per user.	Risk: uncontrolled token burn. Benefit: \$ per accepted line — the unit economics of AI coding, measurable per team.
Model Adoption	Distinct models in use (15 in the example estate) ranked by users — claude-haiku-4-5, claude-opus-4-8, composer-2.5-fast, grok variants, defaults.	Risk: unapproved or deprecated models writing code. Benefit: standardization, policy per model tier, negotiation leverage with vendors.
Agentic Application Discovery	Repositories with AI/LLM usage: external LLM APIs, agentic frameworks (LangChain, LangGraph, baml, MCP, CrewAI), self-hosted open models; providers ranked (Anthropic, AWS Bedrock, open source...).	Risk: unknown AI supply chain in own products — EU AI Act / CRA exposure. Benefit: a machine-readable AI inventory (AI-BOM) per repository.
Repository-Level Detection	Per-repo badges: LLM provider, framework, packages, agentic configuration, hosting class (cloud-managed vs. self-hosted), exact model strings, files, last-detected date.	Risk: self-hosted vs. cloud data-exposure paths invisible to AppSec. Benefit: targeted architecture & data-flow reviews instead of blanket audits.
Agent Identity Governance	Each agent a first-class identity with Insights / Activity / Access tabs — access reviewable like any human account.	Risk: agents outside every IAM process. Benefit: least privilege and access reviews extended to AI — one governance model.
OverPrivilege Score & Blast Radius	Per-agent OverPrivilege Score (e.g. Cursor: 60) plus top human collaborators and most active repositories.	Risk: over-permissioned agents with large blast radius. Benefit: prioritized remediation — fix the riskiest agent access first.

Feature set as shown in BlueFlag AI Insights (Insights → AI, Identity Governance → AI Agents, Agentic Applications); example values from a live demo estate.

THE PLATFORM VIEW

One module. Five lenses on agentic AI.



Agent Activity

Commits, PRs, reviews per agent



Shadow-AI Discovery

Unregistered agents & identities



AI Economics

Seats, spend, tokens, accepted code



AI-BOM

Models, frameworks, providers per repo



Agent Governance

Access, OverPrivilege, blast radius



SDLC Agent Activity Graph

Active agents with real 90-day activity — commits by agentic AI (AI vs. co-authored), PR activity — and every commit attributed **agent-primary vs. human-primary**.

OUTCOME

Provenance for every AI-touched change — the audit trail regulators and reviewers ask for.



Shadow-AI & Agent Discovery

Discovered agents nobody registered, with identity and repository counts — the difference between the AI estate you licensed and the one you actually run.

OUTCOME

Complete inventory first, governance second — no agent changes code unseen.



AI Coding Economics

License utilization (active vs. seats), total spend with per-day trend, **switchable Spend / Lines accepted / Tokens**, and model adoption across the estate.

OUTCOME

Cost per accepted line and per team — seats right-sized, shadow usage formalized, ROI measurable.



Agentic Application Discovery (AI-BOM)

Repositories with LLM usage classified by API type, agentic framework (LangChain, LangGraph, baml, MCP, CrewAI), provider and hosting — down to exact model strings per repo.

OUTCOME

A machine-readable AI inventory — the artifact EU AI Act, CRA and supply-chain reviews assume exists.



Agent Identity Governance & OverPrivilege

Every agent is a first-class identity with Insights, Activity and **Access** tabs, an **OverPrivilege Score**, top human collaborators and most-active repositories — reviewable in the same access-review cycle as any employee.

OUTCOME

Least privilege extended to AI — the riskiest agent access remediated first, blast radius known before an incident.

AGENTIC AI ECONOMICS — THE VALUE CHAIN THE MODULE MAKES MEASURABLE

Spend



Token consumption



Accepted code



Features delivered



Business value

Each link is instrumented: spend and tokens per tool/user · lines accepted per tool · agent-attributed commits & PRs per repository · repository-to-application mapping.

THE RETURN ON AI VISIBILITY

4 active · 0 seats

Claude Code in productive use without a single licensed seat — **shadow usage surfaced** and turned into a procurement decision.

\$ / accepted line

Spend, tokens and lines accepted in one view — **the unit economics** that decide which tools and models earn their cost.

15 models

Model adoption made visible — **consolidation and policy** per model tier instead of uncontrolled sprawl.

Score 60

An agent's OverPrivilege Score plus its most-active repos — **remediation prioritized** by blast radius, not alphabetically.

WHY EACH STAKEHOLDER SHOULD CARE

CISO	Every agent an identity with an access trail and an OverPrivilege Score — Shadow AI eliminated as a category , agent-written code attributable in any incident or audit.
CFO / FinOps	Seats vs. active users, spend vs. tokens vs. accepted code — license right-sizing and \$-per-accepted-line replace gut feeling in AI budget decisions.
CTO / Platform	Which agents, models and frameworks actually deliver — velocity data per tool to standardize the stack without slowing teams down.
IAM leader	Agents join the same access-review cycle as humans — one least-privilege model across employees, contractors, NHIs and AI.
GRC / compliance	AI-BOM per repository and provenance per commit — the evidence EU AI Act, CRA, NIS2 and internal AI policies assume , generated continuously.

EXECUTIVE MESSAGE

"Agentic AI is the fastest-growing workforce in your software factory — and the only one with no badge, no manager and no cost center. BlueFlag gives it all three: an identity, a governance model and a P&L."

BlueFlag Security · Prevent first. Prove always.

See your real agentic AI estate — agents, spend and shadow usage.

Read-only AI Insights scan · deploy in <30 minutes · first findings in 48 hours

blueflagsecurity.com

Informational solution brief. All figures and screenshots referenced originate from a live example environment (90-day window) and are illustrative; customer-specific values are produced by the read-only scan. Feature availability per BlueFlag's current AI Insights module (beta features labeled in-product).