

Flux platform overview

As AI code generation scales across engineering organizations, the tools that analyze and govern that code require the same scrutiny as the code itself.

Flux is a code-first engineering intelligence platform that gives engineering leaders ground-truth visibility into how work actually happens across their GitHub estate. Flux connects directly to your GitHub organization in minutes via the GitHub App integration, analyzing real commits, pull requests, and repository activity. It is read-only, and requires no tickets, no tagging, no changes to how your teams work, and no elevated permissions.

→ Why engineering leaders use Flux

Ground-truth visibility into what's being built, derived from the codebase rather than tickets or self-reported status updates

Clear, defensible answers to board and executive questions about AI coding outcomes, delivery, and risk

→ Integration and deployment

Source integration

Flux connects natively with GitHub Cloud and GitHub Enterprise organizations.

Deployment models

Multi-tenant SaaS (cloud-hosted)

Repository access

Accesses repository activity, including commits, pull requests, repository structure, and metadata, to perform automated code analysis and generate insights. It does not use access tokens that can be lost or leaked. App permissions are reviewable and revocable at any time from your GitHub settings. Flux does not modify or write to your repositories, execute code, or require elevated permissions.

→ Core capabilities



Codebase insights & CapEx/OpEx

Synthesized codebase analysis at team/org level, automated feature/maintenance/bugfix/refactor tagging, CapEx vs OpEx categorization, export insights directly into board decks, and data annotations to track changes, PTO, and external events affecting org performance, and R&D tax credit-ready documentation generated automatically from work type classification.

Data analyzed

Commit history, PR metadata, annotation events, R&D-eligible activity documentation.

Questions you can answer

Where's engineering investment going? What's CapEx vs OpEx this quarter? Can I move this into a board deck without rebuilding it and is it defensible? Can I substantiate our R&D tax credit claim without pulling engineers off their work?



Real-time tracking

Continuous repository tracking showing what engineering teams are actually building, derived directly from the codebase rather than ticket status. Includes deployment frequency, lead time for changes (P50/P90 vs historical baseline), recent activity summaries, and repository-wide activity views.

Data analyzed

Commits, branch activity, pull request lifecycle, deployment events, lead time intervals.

Questions you can answer

Is this feature actually shipping? Where is effort concentrating? What changed in the codebase this week that didn't show up in a ticket? Is our shipping cadence slowing down? How long does work actually take to reach prod?



AI impact analysis

Measures outcomes of AI coding tool adoption and its downstream effects on development workflows, including velocity changes, code churn, and work type distribution.

Data analyzed

AI Assisted Authorship rates, cycle time, work type classification.

Questions you can answer

Is AI adoption changing what we ship, or just how much we commit? What did we actually get for our AI investment?



Ask Flux

Natural-language query interface. Ask questions about PR trends and delivery data directly, get answers without filing a request with a data team.

Data analyzed

PR history, commit activity, delivery metrics already tracked by the platform.

Questions you can answer

What's our PR trend this month? Why did lead time spike last week?

→ Core capabilities, continued



Security and risk alerts

Automated static code scanning to identify structural vulnerabilities, exposed secrets, and technical debt. Security scanning runs on every pull request. Snyk integration surfaces existing security signal alongside delivery data. Tracks introduced dependencies and associated vulnerabilities. Traces issues back to the introducing commit.

Data analyzed

Exposed credentials and secrets, CVE indicators, technical debt hotspots, dependency changes, PR size, reviewer count, and review conversation.

Questions you can answer

What vulnerabilities are in code that hasn't been reviewed yet? Where is technical debt accumulating? Are there exposed secrets in recent commits? What dependencies did we just introduce, and are they safe?



Stability & Reliability

Manually logged incidents + reverts inferred from PR data, change failure rate tracking, change and failure detail pages for blameless postmortems without reconstructing an incident across five tools.

Data analyzed

Incident logs, PR/revert data.

Questions you can answer

What's our MTTR? What % of deploys cause failures? What happened during this incident, and why?



Team dynamics

Surfaces collaboration patterns, review velocity, and contribution distribution to help leaders understand how their teams are actually working. Includes cross-team comparisons, auto-discovered teams (see how work actually flows across the org), and developer participation ratio (PR contribution-to-review patterns).

Data analyzed

PR review turnaround, code distribution, bottleneck indicators, team-level contribution patterns.

Questions you can answer

Who are the review bottlenecks? Where is knowledge concentrated? Is code review keeping pace with output? How does this team compare to others, not just the org average?

→ Security and data governance

Data isolation

Each customer's data is hosted in a separate PostgreSQL schema, enforced at the database level. All customer data is encrypted at rest and in transit.

Production access controls

Developers do not have direct access to production infrastructure. All changes are defined as code, version-controlled, and merged through human-reviewed pull requests. Customer data is never copied to staging environments, remote or local agentic development environments, or developer machines.

LLM governance

Flux uses Google's Gemini foundational models, accessed via Vertex AI on Google Cloud Platform (GCP). Google and Flux may use chat interaction data, not code or repository data, to improve our products.

→ Security and data governance, continued

Data retention

- Customer data retained for up to 30 days after contract end
- System activity logs (PII-stripped) retained for 90 days after contract end
- Automated tooling enforces retention policies
- Secure deletion procedures applied to data no longer needed

Contractual protections

Contractual confidentiality protections for customer source code and data, formalized in a subscription agreement that requires mutual written consent to modify.

Security program

- Assessments based on Cloud Security Alliance Cloud Controls Matrix (CSA CCM), National Institute of Standards and Technology Cybersecurity Framework (NIST CSF), and Open Web Application Security Project Top 10 (OWASP Top 10)
- Automated Static Application Security Testing (SAST) integrated into every pull request pipeline; merges blocked on identified issues
- SOC 2 and ISO 27001 certifications on the roadmap
- Automated dependency monitoring via Dependabot, with updates typically handled within one business day. Libyear tracks dependency age to maintain currency against security updates.

Encryption standards

Data in transit is encrypted with TLS 1.2 minimum (TLS 1.3 negotiated by default), terminated at a GCP HTTPS load balancer under a RESTRICTED SSL policy with forward secrecy; HTTP is redirected to HTTPS. Data at rest is encrypted with AES-256 using Google-managed keys across Cloud SQL (PostgreSQL), Cloud Storage, and Secret Manager.

Incident response

Documented incident response policy covering detection, triage, containment, customer communication, and post-incident review. Customers are notified promptly upon discovery and verification, following GDPR and CCPA. Specific notification SLAs are available by negotiation.

→ Enterprise readiness

Enterprise SSO

User authentication is handled via Auth0 with support for third-party SSO providers. Flux does not manage or store user passwords.

On-premises deployment

Self-hosted deployment option available in addition to multi-tenant SaaS.

Enterprise security reviews

Flux has passed enterprise security reviews conducted by customers in regulated industries, including cybersecurity.

24/7 support

Dedicated global support.

Ready to see Flux in action? [Request a demo](#) or contact us at info@askflux.ai.