

SDLC Security Risk Assessment

ASSESSMENT & FINDINGS

Delivered by Knovation Solutions

What it is: A security audit of the software development lifecycle, covering every identity that can reach your code: developers, contractors, service accounts, and AI coding assistants and autonomous agents. Read-only API connection to source control, CI/CD and registries. Agentless. No code changes, no developer friction.

When you use it: Where scanners already cover the code but nothing covers who and what is shipping it. SAST, SCA and ASPM examine code. The exposure that leads to real incidents lives in identities, and those tools were never built to see it.

Assessment Overview and Capabilities

AI Governance in the SDLC

- Full inventory of every AI coding tool in use, including Copilot, Claude, Cursor and Codex
- Sanctioned against shadow AI, each tool attributed to a named individual
- Every AI agent listed as an identity with a named owner
- Which developers and teams use each agent, and which repositories it has written into

Insider Threat and IP Loss

- Full roster of internal developers, contractors and vendors with access
- Over-privileged identities ranked by actual usage rather than assigned role
- Risky behaviours flagged: bypassed branch protections, unauthorised merges, activity from unexpected locations
- Anomalous repository activity surfaced, including sudden activity after long dormancy

Non-Human Workforce Governance

- Count of hidden service accounts uncovered
- Stale tokens, API keys and long-lived credentials identified
- Dormant and orphaned machine identities flagged
- Non-human identities mapped to owners and the pipelines they serve

SDLC Posture

- Full catalogue of repositories, CI/CD pipelines and build systems
- Exposed secrets and critical vulnerabilities flagged in code and infrastructure as code
- Repositories lacking signed commits, branch protection or mandatory review
- Ownership and maintenance gaps documented

Compliance Visibility

- Separation of duties and least privilege gaps documented
- AI-authored pull requests bypassing human review, flagged
- Findings evaluated against seven default frameworks, including OWASP, CIS and ISO 27001
- Custom frameworks available where a client standard applies
- Findings mapped to evidence auditors accept

Prerequisites

Required

- Read-only API access to source control management, CI/CD and registries
- One point of contact to authorise the connection and review findings
- No agent installation, no code changes, no build pipeline modification

Optional, extends the assessment

- Identity provider connection, for fuller correlation across human and non-human identities
- Multiple repositories and pipelines connected, for estate-wide rather than project-level findings

Timeline

- **Under 30 minutes:** Read-only API connection to source control, CI/CD and registries
- **Discovery:** Every human, non-human and AI identity inventoried into one identity graph with behavioural baselines
- **Analysis:** Toxic combinations, over-privilege, shadow AI and exposure points correlated across identity, toolchain and code
- **48 hours:** Prioritised findings report reviewed with you, ranked by risk and ready to act on

Assessment Outcomes and Findings

What you receive:

- A single attributable inventory of every developer, contractor, AI tool and identity that can touch your code
- Every service account, bot and AI agent mapped and risk-scored alongside human developers
- A compliance-ready evidence base
- A prioritised, risk-ranked action list your team can execute against immediately
- An evidence-backed rationale for closing the SDLC risk that SAST, SCA, ASPM and IAM tools were never built to cover

Findings reported. Each item below is a control the platform evaluates against your connected environment. Coverage varies with environment configuration and the platforms connected.

Identity and Access

- Every identity able to touch code: developers, contractors, vendors, service accounts, bots and tokens
- Dormant identities whose access still resolves
- Dormant identities still holding administrative rights over repositories
- Identities not linked to corporate identity management
- External and guest identities, distinguished from internal
- Tokens and keys that continue to function after the account that issued them has gone

Entitlements and Privilege

- Access granted and never used, split by administrative, write and read
- Access granted directly to an individual rather than through a group
- Teams containing inactive members, and teams holding repository access nobody uses
- Right-sizing opportunities ranked by usage rather than assigned role

Behaviour

- Activity from locations that are rare for the tenant
- Activity at unusual hours or on unusual days
- Activity on repositories after long dormancy
- Commits where the committer and the author do not match
- Branch protection bypassed by administrators
- Repositories cloned, made public or deleted, tied to the identity that did it

Repository, Branch and Pipeline Posture

- Repositories with unprotected default branches
- Protected branches permitting a single approver, or missing code owner approval
- Repositories carrying more administrators than policy allows
- Pipeline hardening gaps: fork protection, token scoping, job authorisation scope
- Pipeline definitions held outside source control, where changes bypass review

Code, Secrets and Dependencies

- Credentials committed to source code, with those confirmed to still authenticate separated from the rest
- Critical and high vulnerabilities in open-source packages
- Infrastructure as code carrying critical or high misconfigurations

AI Coding Tools

- Every AI coding tool in use, discovered through repository artefacts rather than endpoint telemetry
- Each tool attributed to the individuals using it and the repositories it has written into
- Sanctioned tools separated from those in use without approval
- Pull requests approved or merged by an AI agent rather than a person

What the Assessment Does Not Do

- It does not scan code for logic defects or quality issues
- It does not block, gate or halt a build
- It does not modify anything in your environment
- It does not see activity in platforms that are not connected
- Some platforms do not emit the events required for certain detections

Closing

We appreciate the opportunity to support your organisation in strengthening its software development lifecycle security. Please contact us at info@knovation.co.za with any questions regarding this assessment.