GitLab + Sonatype + BlueFlag Security

The Three-Layer SDLC Security Stack

Where They Fit Together

The Scenario

You've invested in GitLab and Sonatype. Your code is scanned. Your dependencies are managed.

But your team still can't answer: who made this change, did they have proper access, is this behaviour normal, and where's the audit trail?

This matrix shows where identity governance fits and why it matters.

What GitLab and Sonatype Do Well

Capability	GitLab	Sonatype	What This Solves
Code scanning (SAST)	Native	Limited	Finding bugs in your code
Dependency scanning (SCA)	Basic/Good	Market leader	Finding vulnerable open-source libraries
Container security	Native	Partial	Finding vulns in container images
Artifact/package management	Basic	Strong (Nexus)	Managing and firewall-ing artifacts
CI/CD pipeline execution	Core	N/A	Building and deploying software reliably
SBOM generation	Basic	Strong	Tracking what's in your software
Pipeline visibility	Strong	N/A	Seeing what code was built, when, where

Together, they answer: what's vulnerable in my software and dependencies, and how is it being built?

The Critical Gap: Identity and Governance

When a GitLab scan flags a vulnerability or Sonatype surfaces a risk, your team has found a technical problem. But without identity governance, they cannot answer the critical questions that follow.

The Three-Tool Stack: Executive View

Question	Answer	Tool
How is software built?	Repositories, pipelines, releases, collaboration	GitLab
What are we using in our software?	Open-source dependencies, artifacts, SBOMs, vulnerability data	Sonatype
Who built it, who approved it, what AI contributed, are we compliant?	Identity, behaviour, governance, evidence, regulatory alignment	**BlueFlag**

Together: GitLab executes. Sonatype secures components. **BlueFlag** governs who, why, and whether the work meets policy.

Critical Questions GitLab + Sonatype Cannot Answer

Question	GitLab	Sonatype	Impact if unanswered
Who made the commit?	Log entry only	No visibility	You know code changed; you don't know if that identity should have access
What was their permission level?	No visibility	No visibility	A contractor with read-only access merging critical code looks the same as a principal engineer
Is this behaviour normal for that identity?	No analysis	No visibility	Dormant account suddenly pushing production code is a breach signal; you won't see it
Who approved the merge?	Git log; human-based	No visibility	No correlation between who changed what and who was supposed to approve it
Does this identity have orphaned or excessive permissions?	No visibility	No visibility	20-40% of identities in most pipelines have standing permissions they shouldn't have
Should this service account have access to production?	No governance	No visibility	Service account with stale PAT commits malicious code; you have no way to track or prevent it
Who deployed this to production?	Log entry	No visibility	You know it was deployed; you don't know if the deploying identity had proper approval
Is there audit evidence linking identity, action, and approval?	Partial	No	SOC 2 and ISO 27001 auditors want: identity X with role Y approved by Z performed action A on date D. You have pieces, not the story

What Happens When These Questions Aren't Answered

Operational Risk

Scenario: Sonatype flags a critical CVE in a dependency used in your payment processing service.

With GitLab + Sonatype only, you know: the vulnerability exists, it's in a component you use, a merge happened at 2:15pm on Tuesday.

You don't know: whether the identity that committed it still has access, whether they're behaving normally, whether they should have had that level of access, whether a service account is also using that library, whether an AI agent contributed to the problematic code.

Result: you can't answer the security buyer's question: who's accountable, and did we have controls?

Compliance Risk

Scenario: your CISO runs a SOC 2 audit. Auditor asks: walk me through a recent change in production. Who made it, with what permissions, who approved it, and where's your audit trail?

With GitLab + Sonatype only, you piece together: git log (who committed), pipeline log (what was built), Sonatype report (what components are used).

You're missing: identity privilege proof at the time of action, behavioural baseline (was this normal?), approver role confirmation, non-human identity audit trail (service accounts, tokens, bots, AI agents).

Result: partial audit evidence. The auditor notes a control gap: identity governance not evident.

Insider Threat & Compromise Risk

Scenario: your CI/CD pipeline has 50 developers, approximately 2,250 non-human identities (service accounts, PATs, tokens, runners, AI agents).

With GitLab + Sonatype only, you have: visibility into code changes, visibility into dependencies.

You don't have: a complete identity graph of who/what touches the pipeline, behavioural analytics on identity actions, detection of dormant accounts reactivated, detection of excessive privilege usage, detection of abnormal patterns (bulk secret access, off-hours commits, etc).

Result: a compromised service account pushing malicious code is invisible until forensics, not prevention.

Where BlueFlag Fits: The Identity Governance Layer

BlueFlag operates at a different layer than GitLab and Sonatype. It doesn't replace them; it completes them.

Full Identity Graph Maps every developer, contractor, service account, PAT, token, bot, and AI agent across GitLab, Sonatype, cloud IAM, and external systems. Every identity tied to privilege level, behaviour baseline, and approval chain.

Behavioural Risk Correlation Correlates GitLab pipeline signals (commits, merges, deploys) with identity behaviour patterns to detect anomalies in real time. Surfaces dormant accounts reactivating, unusual access patterns, and off-hours activity.

Non-Human Identity Governance Discovers and governs all service accounts, PATs, tokens, bots, and AI agents. Identifies orphaned accounts, excessive permissions, and stale credentials.

AI Coding Agent Governance Auto-discovers which AI coding agents, including Claude Code, Copilot, Cursor, Codex, and GitLab Duo, are in use. Tracks code contribution by agent and detects AI-only commits that bypass review. Enforces policy in real time (block unauthorised agents, enforce approved AI only).

Risk Contextualisation Contextualises findings with identity data. Same CVE flagged by Sonatype is viewed through the lens of who introduced it, their permission level, their behaviour baseline, and exposure scope.

Compliance & Audit Generates audit-ready evidence continuously. Every identity action linked to governance decision, approval, and compliance framework (SOC 2, ISO 27001, DORA, NIS2, POPIA, JS2 FSP, NIST, EU AI Act).

Capability Comparison: With and Without BlueFlag

Capability	GitLab + Sonatype	+ BlueFlag
Find vulnerable code	Yes, full	Yes, full (same)
Find vulnerable dependencies	Yes, full	Yes, full (same)
Know who made the change	Partial, log entry only	Yes, identity graph, behaviour, permissions
Know if change maker had proper access	No	Yes, with permission context
Detect abnormal identity behaviour	No	Yes, real-time alerts
Govern non-human identities	No	Yes, full discovery and governance
Govern AI agents	No	Yes, full governance and traceability
Provide compliance audit evidence	Partial	Yes, full identity governance trail
Show who's accountable for a change	Partial, git log only	Yes, identity + behaviour + approval + audit
Detect and respond to compromise faster	No	Yes, early anomaly detection enabling faster response
Executive/board reporting on SDLC risk	Partial, technical only	Yes, risk-based governance
Regulatory mapping (NIST, ISO, DORA, POPIA)	Partial	Yes, full governance framework (NIST, ISO 27001, DORA, NIS2, POPIA, JS2 FSP, EU AI Act)

Entry Point: Free SDLC Risk Assessment

We offer a 48-hour assessment of your GitLab and Sonatype environment. We map all identities, flag orphaned accounts, excessive permissions, and behaviour anomalies.

Most organisations discover material findings: typically 20-40% of identities are orphaned or over-provisioned.

This assessment gives you concrete data on the identity governance gap and helps you prioritise **BlueFlag** on your roadmap.