

AI-Driven Application Security with Context Intelligence Graph from Cyscale

January 19, 2026

By: [Katie Norton](#)

IDC'S QUICK TAKE

Cyscale announced the Context Intelligence Graph, expanding its platform to formalize the relationship between security signals, decisions, and their outcomes over time. The approach addresses a critical gap as security moves toward autonomous agents: most platforms record what was found, but few preserve why teams acted the way they did.

PRODUCT ANNOUNCEMENT HIGHLIGHTS

On January 21, 2026, Cyscale [announced the Context Intelligence Graph](#), an expansion of the graph-based architecture underpinning its platform. The updated model extends the prior Risk Intelligence Graph to incorporate application security testing, software supply chain security, and application security posture management (ASPM) data into a shared context layer. According to Cyscale, the architecture is designed to improve how security signals are related, interpreted, and retained across the software development lifecycle.

The Context Intelligence Graph is structured around what Cyscale refers to as the five-dimension system of context. This framework organizes application security context across five core dimensions: temporal context (when changes occurred), causal context (how events are related), semantic context (what findings and signals mean), attribution context (who owns or approved changes), and outcome context (what happened after a decision was made). These dimensions are represented through different types of relationships, or “joins,” within the graph, enabling correlations across code, pipelines, artifacts, cloud infrastructure, identities, and runtime environments.

Cyscale also introduced decision traces as an output of the Context Intelligence Graph. These are persistent, queryable records that capture the security decision taken (e.g., prioritizing a vulnerability, approving an exception, or triggering remediation) along with evidence, constraints, rationale, and subsequent outcomes associated with that decision. According to Cyscale, decision traces are intended to function as institutional memory for application security teams and as structured input for AI-driven workflows.

Cyscale positioned the Context Intelligence Graph as a foundation for planned AI capabilities within the platform, including more autonomous and agent-driven security actions. No specific new AI agents or autonomous workflows were launched as part of this announcement. The company stated that the updated graph architecture is designed to support explainability, learning from historical outcomes, and human oversight as AI-driven features are developed.

IDC'S POINT OF VIEW

Cyscale's announcement addresses a period in which decision quality has become a limiting factor in application security. Over the past decade, tooling has expanded through additional scanners, coverage domains, and detection capabilities. Most organizations now face

substantial finding volume. The bottleneck has shifted from detection to interpretation: determining which findings warrant immediate remediation, which can be deferred, and which represent acceptable risk.

This decision challenge is compounded by how application risk behaves in modern software environments. Risk evolves as code moves through pipelines, incorporates dependencies, deploys to infrastructure, and operates with specific permissions at runtime. A vulnerability identified in source code may never reach production, or it may be exposed on internet-facing infrastructure with elevated privileges. The significance of a finding depends on its relationship to these broader system changes, not on the finding itself.

This need for relational visibility intensifies as AI-driven security capabilities are being adopted across application security programs. IDC's *Agentic Application Development and DevOps Survey* (December 2025) found 57% of organizations globally currently use agentic AI tools for application security activities, with another 34% actively exploring adoption. AI systems lacking application context default to generic recommendations that ignore deployment realities and organizational policies. When AI operates autonomously, incomplete context stops being an information quality problem and becomes an execution problem: wrong decisions happen automatically, at scale, and without intervention.

A notable aspect of Cycode's announcement is the use of the Context Intelligence Graph to model how security teams make decisions. The five-dimension system of context captures factors that practitioners consider in practice — sequence, causality, meaning, ownership, and downstream impact — but that have often remained implicit or inconsistently recorded. By encoding these dimensions into a data model, Cycode aims to make decision logic explicit and durable, allowing both humans and AI systems to reference, audit, and reason over past security decisions.

Decision traces extend this model by preserving the reasoning behind security actions over time. Many security platforms capture metadata about findings, yet the rationale for prioritization, exception handling, or deferred remediation is often dispersed across tickets, chat tools, or individual human memory. Without this continuity, AI-driven systems must repeatedly relearn preferences and decision criteria, increasing the likelihood of recommendations that conflict with historical outcomes or accepted risk tolerance.

Although Cycode positions the Context Intelligence Graph as a foundation for AI-driven security decisions within its platform, the architectural implications can potentially extend beyond internal workflows. As AI coding assistants and agents become more integrated into software creation, their ability to generate secure code depends on access to an authoritative security context rather than abstract best practices. Architectures that preserve decision rationale, constraints, and outcomes as a system of record become relevant not only for application security programs, but also for informing AI agents earlier in the development lifecycle.

This announcement from Cycode emphasizes context modeling over feature-level AI augmentation. Many vendors have focused on delivering AI-assisted remediation, conversational interfaces, or automated triage layered on top of existing data models. While such capabilities can improve efficiency, isolated AI features, even when effective individually,

lack the continuity and memory required for sustained, and potentially autonomous, decision-making.

The effectiveness of a context-driven architecture, however, also depends on the quality of underlying data. Maintaining a high-quality context graph depends on the completeness and reliability of underlying data sources. Platforms that rely heavily on third-party integrations, as is common in ASPM-oriented architectures, have limited control over how upstream data is generated, normalized, and refreshed. These dependencies limit the degree to which vendors can validate or correct ambiguous signals, reinforcing the risk that weak inputs undermine downstream reasoning.

Platforms that operate native scanners, like Cycode, may have a greater ability to validate or enrich certain data streams, even when those tools are not the primary source of detection. This refinement does not eliminate data quality challenges, but it can expand a vendor's influence over how signals are interpreted and connected. As security platforms move closer to autonomous operation, deficiencies in input quality do not merely reduce insight; they propagate assumptions directly into automated actions, amplifying the effects of "garbage in, garbage out."

In addition to supporting autonomous action, capturing decision rationale, encoding constraints, and linking actions to outcomes enable auditability, oversight, and consistency with organizational risk tolerance. The importance of these capabilities aligns with IDC's prediction that by 2029, growing risks and complexity in multiagent orchestration will drive a 35% increase in investment focused on AI governance and oversight (see *IDC FutureScape: Worldwide Developer and DevOps 2026 Predictions*, IDC #US53858525, October 2025). Architectures that incorporate governance into the decision model are better positioned to support autonomy while maintaining control mechanisms.

For security leaders, this announcement represents both a development to evaluate within Cycode's roadmap and an indicator of market direction. As "context" becomes an increasingly common, and often diluted, term, Cycode is positioning it as a structured record of how security decisions are made, constrained, and evaluated over time. This decision-centric approach moves beyond generic correlation layers and positions the Context Intelligence Graph as a reasoning layer intended to support accountable autonomy. As Cycode develops agentic capabilities on this foundation, its ability to demonstrate measurable improvements in prioritization accuracy, remediation outcomes, and operational consistency will be central to validating the architectural model.

Subscriptions Covered:

[DevSecOps, Vulnerability Management, and Software Supply Chain Security](#)

Please contact the IDC Hotline at 800.343.4952, ext.7988 (or +1.508.988.7988) or sales@idc.com for information on applying the price of this document toward the purchase of an IDC or Industry Insights service or for information on additional copies or Web rights. Visit us on the Web at www.idc.com. To view a list of IDC offices worldwide, visit www.idc.com/offices. Copyright 2026 IDC. Reproduction is forbidden unless authorized. All rights reserved.