IDC Link
Real-Time IDC Research® Opinion on Industry News, Trends, and Events

F5 Acquires CalypsoAI: GenAI Red Teaming Is Officially a Platform Feature

September 11, 2025

By: Frank Dickson, Grace Trinidad

IDC'S QUICK TAKE

F5 is acquiring CalypsoAI to integrate GenAI red teaming, threat detection, and governance into its Application Delivery and Security Platform (ADSP). This move reflects the growing demand for scalable AI security as regulations like the EU AI Act gain traction. CalypsoAI's tools will enhance F5's capabilities across models and environments. The AI red teaming market is rapidly evolving, with major players like Cisco, Palo Alto Networks, and SentinelOne also making acquisitions. Other vendors are building native GenAI protections. The acquisition signals that GenAI red teaming is now a core platform feature, not a standalone product category.

M&A ANNOUNCEMENT HIGHLIGHTS

F5 has announced its intent to acquire CalypsoAI in a move that brings runtime protection, threat detection, and governance for AI systems into F5's ADSP. As organizations increasingly embed AI into their operations, the traditional security stack has struggled to keep pace with emerging risks like prompt injection and shadow AI. CalypsoAI's capabilities aim to fill that gap for F5, offering adaptive guardrails and visibility across models, clouds, and environments, resulting in end-to-end protection for any AI model from pilot to production, as stated by Francois Locoh-Donou.

The acquisition reflects growing demand for scalable, model-agnostic AI security solutions that keep pace with a rapidly evolving AI landscape, especially as regulatory frameworks like the EU AI Act gain traction. CalypsoAI's red teaming and observability tools will complement F5's existing strengths in app and API delivery, creating a unified platform for securing the full AI stack. Financial terms weren't disclosed, but the deal is expected to close by the end of F5's fiscal fourth quarter.

IDC'S POINT OF VIEW

The AI red teaming market emerged early in 2025 as a specialized segment within application security, focused on proactively identifying and mitigating vulnerabilities in AI systems, particularly those powered by generative models. IDC discussed the segment in its "IDC Innovators: AI Red Team Platforms, 2025" document in April 2025. These platforms simulate adversarial attacks, assess model robustness, and monitor runtime behavior to ensure AI systems remain secure as they evolve. Key capabilities include

adversarial testing (e.g., prompt injection, data poisoning), continuous monitoring, and compliance reporting.

Interestingly, the market is maturing faster than expected. IDC notes that AI red teaming is being absorbed into broader cybersecurity platforms before fully developing as a standalone category. The result? A race to integrate red teaming as a feature rather than a product, with acquisitions leading the charge. Cisco kicked things off by acquiring Robust Intelligence, a move that brought AI model validation and runtime protection into its security stack. Palo Alto Networks followed suit with its acquisition of Protect AI, a mature player offering model scanning, posture management, and red teaming capabilities — now fueling the development of Prisma AIRS. Most recently, SentinelOne acquired Prompt Security, a Tel Aviv–based start-up focused on securing GenAI applications from threats like prompt injection and shadow AI. Now, F5 is acquiring CalypsoAI. With these moves, the market is starting to resemble a Marvel crossover event — everyone's assembling their AI security Avengers.

While some vendors have opted for acquisitions to enter the space, others are building their capabilities from the ground up — proving that not every cybersecurity hero needs a merger montage.

Cloudflare has introduced AI Prompt Protection, a feature within its Cloudflare One suite that acts as a real-time filter for GenAI interactions. It monitors prompts and responses across platforms like ChatGPT, Claude, and Gemini, detecting risks such as data leakage, toxic content, and shadow AI usage. The system enforces enterprise policies, blocks unsafe prompts, and even redacts sensitive information — like a digital bouncer checking IDs and confiscating contraband at the door. It's designed to be frictionless for users while giving security teams visibility and control over AI usage.

Akamai launched its Firewall for AI to protect chatbots, copilots, and LLMs from threats like prompt injection, jailbreaks, hallucinations, and data exfiltration. The firewall supports edge and cloud-native deployments and integrates with Akamai's broader security stack, including WAAP and bot management. It's model agnostic and compatible with platforms like Hugging Face, LangChain, and OpenAI. Beyond blocking malicious inputs and outputs, it offers advanced content moderation and compliance features — think of it as a multilingual AI bodyguard with a law degree and a strong moral compass.

Trend Micro has leaned into its research roots, expanding its Zero Day Initiative to include GenAI vulnerabilities. It's actively red teaming inference servers, vector databases, and containerized AI environments. Trend's approach is comprehensive, targeting the full AI stack and emphasizing proactive threat discovery.

CrowdStrike has embedded GenAI security into its Falcon platform, offering AI Red Team Services and adversarial testing aligned with OWASP standards. Its Charlotte AI enhances detection and response, making GenAI security a native feature rather than an add-on.

To state the point overtly, the F5 acquisition of CalypsoAI confirms that GenAI red teaming is officially a platform feature. Stand-alone solutions are workable when there is a lack of solutions. However, platform vendors are added aggressively adding the functionality before GenAI red teaming had the opportunity to become a market. F5 is also positioning its product as future proof in its ability to “secure every app, every API — and now every AI model and agent — anywhere.” This is a notable feature in the vendor's messaging approach that might address instances of AI hesitancy. By reassuring its customers that they will remain secure even as the AI market and use cases evolve, they enable customers to move forward with their AI and agentic adoption with some confidence that the AI security solution is prepared to meet any changes. New protocols? Co-pilots? Agents? Swarm intelligence? A year in AI is comparable to a dog year. Security and data risks remain the most common inhibitors to AI adoption. Assuaging these concerns for the long term is the prime objective.

Subscriptions Covered:
Trust Measurement and Metrics