451 Research
by S&P Global

Market Insight Report Reprint

With Continuous Red Teaming, ASAPP aims to set a new standard for AI trust

August 11, 2026

by Raul Castanon-Martinez

The provider is betting that enterprises will now buy AI based on verifiable safety, not just capability. Its new Continuous Red Teaming offers a self-hardening architecture that automatically finds and fixes weaknesses, providing the auditable security metrics that static approaches lack.

Introduction

In May, agentic customer experience platform provider ASAPP announced the launch of Continuous Red Teaming — a new capability that integrates adversarial AI testing directly into ASAPP's model evaluation framework. Built on Promptfoo, an AI security platform recently acquired by OpenAI, the capability continuously screens for more than 50 vulnerability types, giving enterprise customers real-time data to trust their AI in production. The announcement marks a meaningful step forward for the company, building on its early positioning with GenerativeAgent as a differentiated approach to agentic customer service. Since then, ASAPP has sharpened its identity around measurable, production-grade AI safety — a differentiator that resonates with enterprise customer experience and IT security buyers navigating an increasingly complex agentic AI landscape.

THE TAKE

ASAPP's move to continuous, automated red teaming addresses a genuine gap in enterprise AI governance: the transition from static, periodic security audits to always-on adversarial validation. By framing safety as a quantified, auditable posture — anchored by attack success rate metrics and aligned to OWASP and NIST frameworks — ASAPP is positioning itself as the trust layer for enterprises deploying AI in high-stakes, customer-facing environments. The integration with Promptfoo and the outcome-based pricing model further reinforce ASAPP's enterprise credentials. While challenges remain in competing with larger CCaaS incumbents, ASAPP's security-first architecture and HILA governance model create a defensible differentiator for both chief information security officers and customer experience leaders.

Context

The rapid proliferation of agentic AI in the enterprise is creating a new and complex security paradigm. These systems introduce novel attack surfaces — including indirect prompt injections via poisoned data, tool misuse that can enable lateral network movement and cross-session PII leakage — that static safety filters and traditional security models were not designed to address.

This expanding threat landscape is no longer just a technical concern; it is rapidly becoming a core compliance issue. Regulatory pressure is intensifying, with AI red teaming increasingly required to demonstrate compliance with critical frameworks like the NIST AI Risk Management Framework (RMF) and the EU AI Act. These regulations are fundamentally shifting AI safety from a voluntary best practice to a legal and operational necessity. To gain or maintain market access, enterprises must now provide verifiable proof of model safety and robustness. This is forcing a market-wide move away from periodic, check-the-box audits and toward continuous adversarial testing as a standard, non-negotiable component of enterprise AI governance.

According to 451 Research's Voice of the Enterprise: AI & Machine Learning, Use Cases 2026 survey, the largest proportion of generative AI initiatives remain in development or confined to departmental pilots. Survey data shows that data privacy, security risks and regulatory compliance issues are the top challenges affecting the adoption or performance of AI initiatives.

Data and governance issues contribute to clogged AI project pipelines

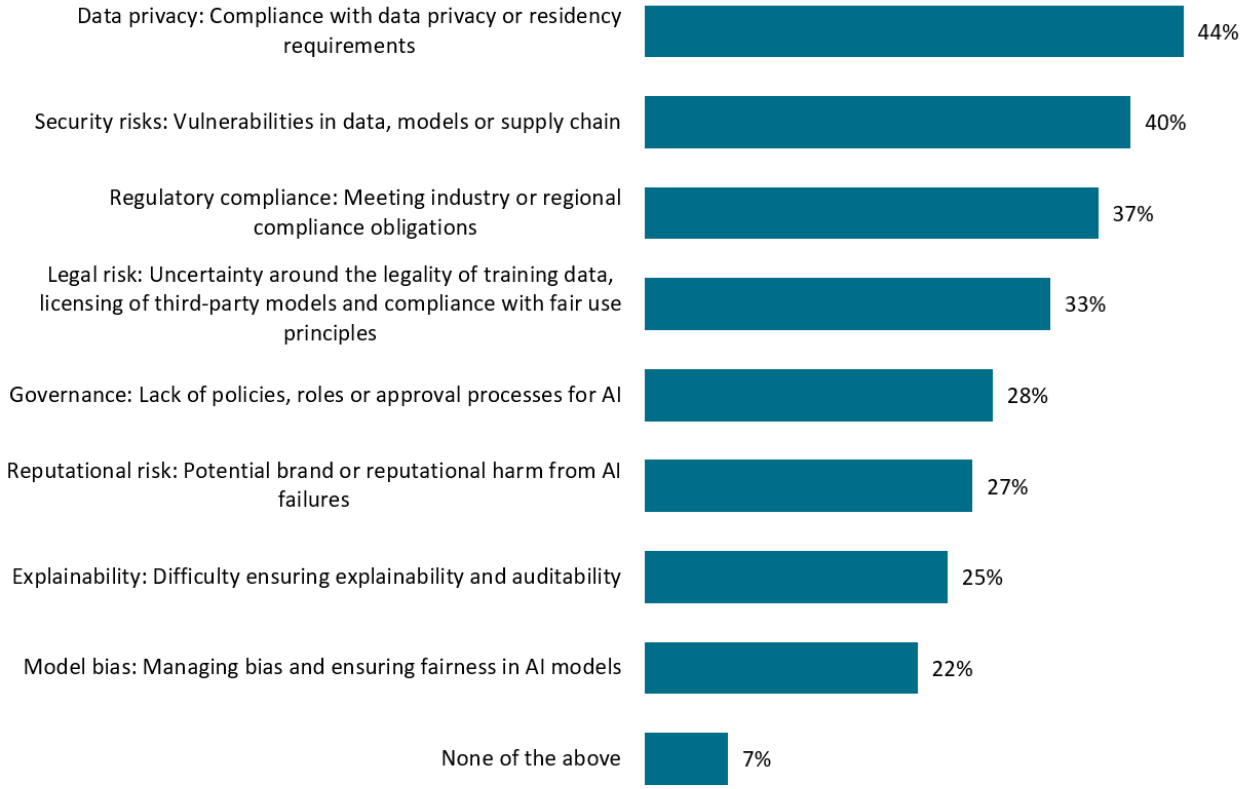

Source: 451 Research by S&P Global.

Voice of the Enterprise: AI & Machine Learning, Use Cases 2026.

Q. Which of the following governance, risk and compliance challenges over the past 12 months impacted the adoption or performance of AI initiatives at your organization? Please select all that apply. Base: All respondents (n=979).

Technology

ASAPP's Continuous Red Teaming capability is powered by Promptfoo's Red Team Agents — high-reasoning, uncensored models that simulate attackers iteratively probing for weaknesses in ASAPP's specific business logic. Rather than relying on static blocked-word lists, the system runs thousands of evaluations across more than 50 vulnerability classes organized around the following three critical security pillars.

Core model integrity & control: Tests for adversarial jailbreaking (including many-shot attacks and character-level obfuscations such as Leetspeak and Base64 encoding), system override and logic hijacking attempts, and hallucinated authority — where a model is coerced into making unauthorized legal or business commitments.

Data privacy & knowledge base security: In retrieval-augmented generation-driven deployments, ASAPP tests for indirect prompt injection via poisoned documents, knowledge base exfiltration and cross-session PII leakage across tenants.

Agentic & operational security: Tests against broken object level authorization (BOLA), server-side request forgery (SSRF) and tool-calling exploitation that could expose underlying system architecture.

The risk benchmarking framework tracks the attack success rate (ASR) for each model update, uses automated high-reasoning model graders to remove human bias, and maps results directly to the OWASP Top 10 for large language models and the NIST AI Risk Management Framework — ensuring safety metrics are audit-ready from day one.

Products

At the core of ASAPP's customer experience platform is GenerativeAgent, an agentic AI designed to autonomously listen, reason, act and improve through interaction intelligence across both voice and digital channels. The agent can resolve customer issues end-to-end or invoke the human-in-the-loop agent (HILA) model — a design philosophy that inverts the traditional escalation model. Rather than AI assisting a human agent, HILA positions the human as a consultant to the AI. The AI handles the interaction and requests human approval or guidance (for example, authorizing a discount) when edge cases arise. This allows enterprises to scale AI-driven service while preserving human judgment for high-risk moments.

For hard escalations, ASAPP provides its own agent desk or integrates with major CCaaS and CRM platforms. Customers may access Continuous Red Teaming, which is already generally available, as part of the ASAPP CXP (Customer Experience Platform) platform fee.

Strategy

ASAPP's commercial model is structured in two components: a platform fee (self-service tier based on the customer's knowledge level and needs) and a variable fee with multiple options — the default being outcome-based, interaction-level pricing. This pricing architecture aligns vendor incentives directly with customer outcomes, distinguishing ASAPP from seat- or usage-volume-based licensing common among legacy CCaaS providers.

The strategic thesis behind Continuous Red Teaming is that safety is a competitive moat, not a compliance checkbox. By automating red teaming into every model iteration cycle, ASAPP can ship new features and model updates faster than competitors that rely on quarterly security reviews. Any newly discovered attack prompt can be immediately templated, enabling ASAPP to test not just that single attack but the entire attack category — a compounding defensive advantage over time.

Further reinforcing this defensive moat, ASAPP uses its Continuous Red Teaming capability internally. This internal use serves as both a powerful signal of confidence and a key customer reference point. The company contends that for enterprise AI, trust must be earned through verifiable evidence, not vendor claims, by providing CISOs and procurement teams with the specific metrics, clear methodology and regular reporting they require. This internal validation is a core component of its strategy to create a safety architecture that automatically identifies and fixes weaknesses.

Competition

ASAPP competes in a dynamic and well-funded CX AI platform market against significantly larger incumbents. These include Genesys Cloud Services, which in March announced its Genesys Cloud platform closed the fourth quarter of its 2026 fiscal year (Nov. 1, 2025–Jan. 31, 2026) with nearly $2.6 billion in annual recurring revenue (ARR), growing more than 35% year-over-year. Meanwhile, in August, NICE announced results for the second quarter ended June 30, 2026, with AI ARR reaching $362 million and now representing 15% of the company's cloud revenue. Five9 Inc. reported revenue for the first quarter of 2026 of $305.3 million, a 9% year-over-year increase with AI revenue specifically accelerating 68% to an annual run rate exceeding $125 million. Another key competitor is Verint, which launched three new agentic AI tools — Workforce Intelligence, Desktop Intelligence and Quality Intelligence — at its Engage 2026 conference, built on its Agent Factory orchestration environment.

None of these competitors have publicly announced an equivalent to ASAPP's continuous, automated red teaming posture or ASR-based risk benchmarking framework — creating a defensible near-term differentiation window, particularly for security-sensitive verticals such as financial services, healthcare and telecommunications.

SWOT Analysis

STRENGTHS

ASAPP's strategic strength is its security-first differentiation in the CX AI market. It claims a first-mover advantage with continuous automated red teaming, providing quantifiable safety evidence via its ASR metric. This security posture is supported by its purpose-built GenerativeAgent and HILA model for human oversight. Commercially, its outcome-based pricing aligns vendor-customer incentives, while alignment with NIST and OWASP standards eases compliance for regulated industries.

WEAKNESSES

The company's primary weaknesses relate to its market scale and maturity. ASAPP has a smaller brand footprint than long-standing CCaaS incumbents Genesys and NICE, and its go-to-market reach is limited. Furthermore, its innovative pricing model requires significant buyer education, potentially creating friction in procurement cycles.

OPPORTUNITIES

ASAPP is poised to benefit from market and regulatory tailwinds. Growing demand for AI security validation, particularly in regulated verticals, creates a direct need for its offerings, a trend amplified by new mandates like the NIST AI RMF and EU AI Act. The company can fill a critical gap as agentic AI adoption outpaces enterprise security frameworks, allowing its security-centric narrative to resonate with the growing influence of CISOs in technology procurement.

THREATS

Threats are mounting from both intensifying competition and execution risk. Incumbent CCaaS vendors are investing heavily in their own AI safety capabilities, which could erode ASAPP's core differentiator. The competitive landscape also faces uncertainty from ecosystem shifts, such as OpenAI's acquisition of Promptfoo. Internally, the model is threatened by rapidly evolving adversarial techniques and the risk that its own pricing complexity could slow enterprise deal velocity.

Source: 451 Research by S&P Global.

The author(s) used a proprietary S&P Global AI platform in the production of this report. It was subsequently peer-reviewed, fact-checked and edited before publication. S&P Global uses generative AI to create content in accordance with our Terms of Use.

CONTACTS

www.spglobal.com
www.spglobal.com/en/enterprise/about/contact-us.html

These materials, including any software, data, processing technology, index data, ratings, credit-related analysis, research, model, software or other application or output described herein, or any part thereof (collectively the "**Property**") constitute the proprietary and confidential information of S&P Global Inc its affiliates (each and together "**S&P Global**") and/or its third party provider licensors. S&P Global on behalf of itself and its third-party licensors reserves all rights in and to the Property. These materials have been prepared solely for information purposes based upon information generally available to the public and from sources believed to be reliable.

Any copying, reproduction, reverse-engineering, modification, distribution, transmission or disclosure of the Property, in any form or by any means, is strictly prohibited without the prior written consent of S&P Global. The Property shall not be used for any unauthorized or unlawful purposes. S&P Global's opinions, statements, estimates, projections, quotes and credit-related and other analyses are statements of opinion as of the date they are expressed and not statements of fact or recommendations to purchase, hold, or sell any securities or to make any investment decisions, and do not address the suitability of any security, and there is no obligation on S&P Global to update the foregoing or any other element of the Property. S&P Global may provide index data. Direct investment in an index is not possible. Exposure to an asset class represented by an index is available through investable instruments based on that index. The Property and its composition and content are subject to change without notice.

THE PROPERTY IS PROVIDED ON AN "AS IS" BASIS. NEITHER S&P GLOBAL NOR ANY THIRD PARTY PROVIDERS (TOGETHER, "**S&P GLOBAL PARTIES**") MAKE ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, FREEDOM FROM BUGS, SOFTWARE ERRORS OR DEFECTS, THAT THE PROPERTY'S FUNCTIONING WILL BE UNINTERRUPTED OR THAT THE PROPERTY WILL OPERATE IN ANY SOFTWARE OR HARDWARE CONFIGURATION, NOR ANY WARRANTIES, EXPRESS OR IMPLIED, AS TO ITS ACCURACY, AVAILABILITY, COMPLETENESS OR TIMELINESS, OR TO THE RESULTS TO BE OBTAINED FROM THE USE OF THE PROPERTY. S&P GLOBAL PARTIES SHALL NOT IN ANY WAY BE LIABLE TO ANY RECIPIENT FOR ANY INACCURACIES, ERRORS OR OMISSIONS REGARDLESS OF THE CAUSE. Without limiting the foregoing, S&P Global Parties shall have no liability whatsoever to any recipient, whether in contract, in tort (including negligence), under warranty, under statute or otherwise, in respect of any loss or damage suffered by any recipient as a result of or in connection with the Property, or any course of action determined, by it or any third party, whether or not based on or relating to the Property. In no event shall S&P Global be liable to any party for any direct, indirect, incidental, exemplary, compensatory, punitive, special or consequential damages, costs, expenses, legal fees or losses (including without limitation lost income or lost profits and opportunity costs or losses caused by negligence) in connection with any use of the Property even if advised of the possibility of such damages. The Property should not be relied on and is not a substitute for the skill, judgment and experience of the user, its management, employees, advisors and/or clients when making investment and other business decisions.

The S&P Global logo is a registered trademark of S&P Global, and the trademarks of S&P Global used within this document or materials are protected by international laws. Any other names may be trademarks of their respective owners.

The inclusion of a link to an external website by S&P Global should not be understood to be an endorsement of that website or the website's owners (or their products/services). S&P Global is not responsible for either the content or output of external websites. S&P Global keeps certain activities of its divisions separate from each other in order to preserve the independence and objectivity of their respective activities. As a result, certain divisions of S&P Global may have information that is not available to other S&P Global divisions. S&P Global has established policies and procedures to maintain the confidentiality of certain nonpublic information received in connection with each analytical process. S&P Global may receive compensation for its ratings and certain analyses, normally from issuers or underwriters of securities or from obligors. S&P Global reserves the right to disseminate its opinions and analyses. S&P Global Ratings' public ratings and analyses are made available on its sites, www.spglobal.com/ratings (free of charge) and www.capitaliq.com (subscription), and may be distributed through other means, including via S&P Global publications and third party redistributors.

S&P Global uses generative AI to create content in accordance with our Terms of Use (https://www.spglobal.com/en/terms-of-use).