SecurityPal AI

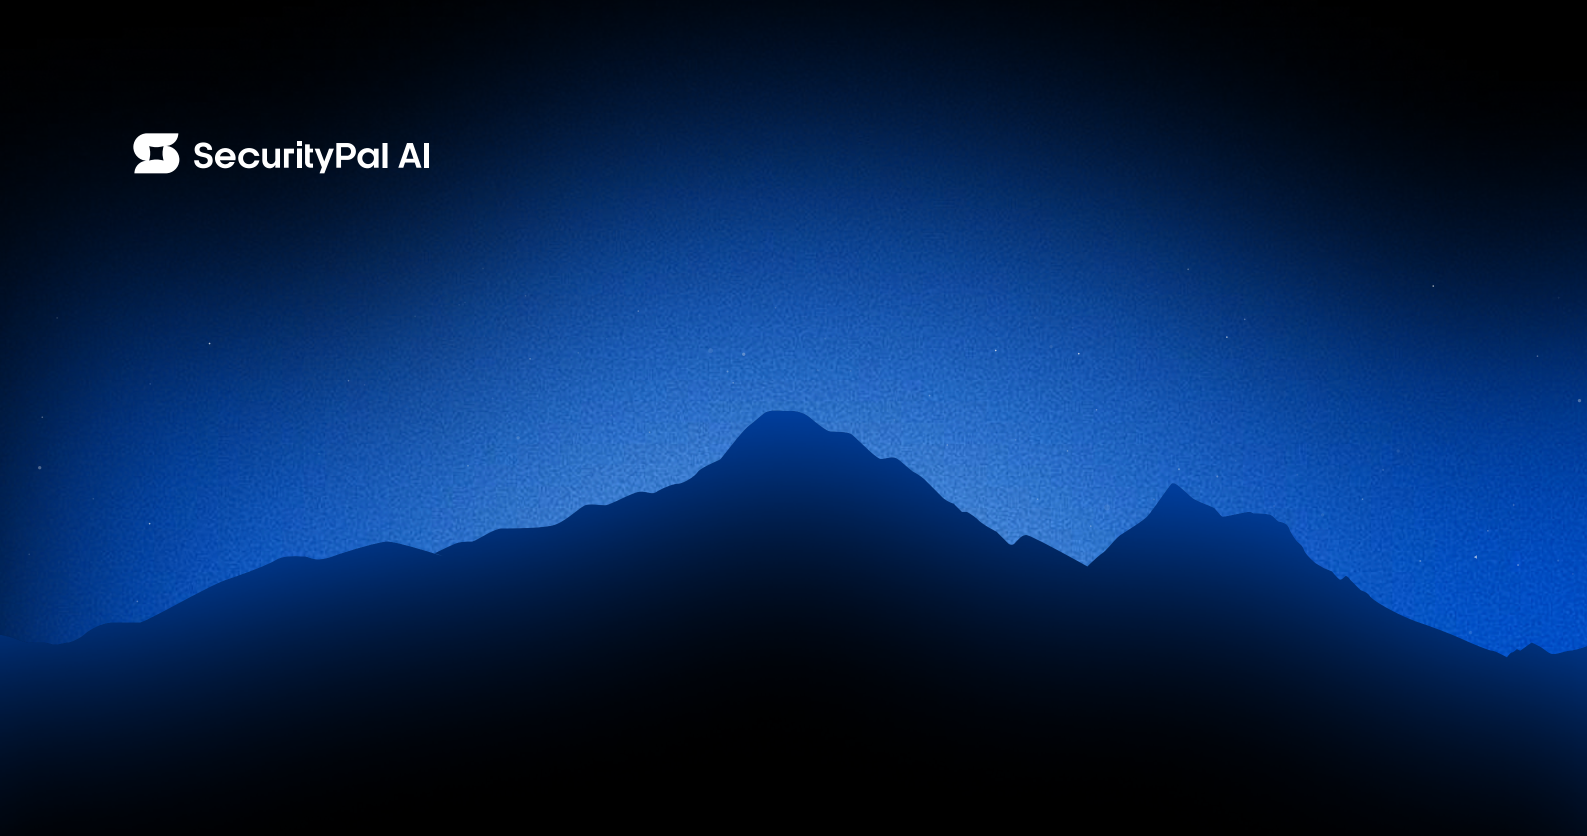

2026 Assurance Insights Report

How Security, GRC, and Trust Are Being Redefined in the Age of AI

✉ contact@securitypalhq.com 📍 Salesforce Tower, San Francisco, CA 94105 ☎ +1 650-842-0352

Executive Summary: Trust at Machine Speed

In 2025, AI completely changed how enterprises build, prove, and retain trust.

AI adoption accelerated across every layer of the enterprise thus powering new products, copilots, internal tools, and customer-facing applications faster and more easily than ever before. At the same time, buyer expectations around security and trust moved just as fast, demanding transparency and governance around emerging AI usage.

This means manual security review processes simply can’t keep up while AI-only assurance management tools can’t capture the nuance required for sensitive security and GRC requests.

2025 marked another year in the ongoing evolution of security reviews shifting from static, reactive checkbox exercises into something far more strategic: a core go-to-market capability that directly impacts deal velocity, buyer confidence, and long-term trust.

SecurityPal offers a uniquely data-driven perspective on how assurance is changing and what leaders must do next. This report is grounded in real data from **thousands of real-world security reviews** providing unmatched visibility into how buyers evaluate trust today.

Key takeaways at a glance:

- **Year-over-year growth** in inbound questionnaires
- Copilot usage more than **doubled** year over year
- Buyer focus shifted from “**AI opt-out**” to “**data-for-training opt-out**”
- **Sharp rise in scrutiny** around AI governance, business continuity, and regulatory readiness
- AI is accelerating **software creation** and dramatically expanding third-, fourth-, and nth-party risk

The Changing Nature of Security & GRC Reviews

From Security Questionnaires to Pre-Sales RFPs

Security reviews are no longer confined to late-stage vendor validation.

What the **data shows:**

In **2025**, we saw a growing volume of questionnaires that:

- Emphasize **product functionality and architecture**, not just security controls
- Appear earlier in the **buying journey**, often before vendors are shortlisted
- Resemble **pre-sales RFPs**, blending security, privacy, and product differentiation

Security teams are being pulled directly into the revenue engine. Assurance is no longer just about mitigating risk. It is increasingly about enabling growth. Security answers are now part of product storytelling, shaping how buyers evaluate credibility and readiness.

What it signals for **2026**:

- Security and GRC teams **must support sales** earlier without burning analyst time
- Assurance must be f**ast, accurate, and repeatable at scale**
- **Trust** is becoming a frontline competitive differentiator

The Rise of **AI Governance** as a **Primary Risk Domain**

In 2025, the rapid advancement and adoption of AI shifted perception. Once seen as a novelty that was challenging to regulate, now became a non-negotiable area of scrutiny.

What's **new** in questionnaires:

We are seeing fewer high-level **"Do you use AI?"** questions. Those are replaced by deep, technical, and governance-focused scrutiny around:

- Model control points (**MCPs**)
- **Human-in-the-loop** workflows
- Use of **customer data** for training
- **Model transparency**, explainability, and oversight

> Buyers **no longer** asked whether vendors use AI. They asked how AI is governed, constrained, audited, and supervised.

2024

“Can we opt out of AI?”

→

2025

“Can you prove our data never trains your models?”

This shift reflects a broader reality: AI is now embedded everywhere, often faster than policies, processes, and controls can mature.

What it signals for **2026**:

AI is making it easier than ever to build and deploy new software, internal tools, and integrations. That velocity is expanding enterprise tech stacks and with them, third-, fourth-, and fifth-party risk.

AI doesn’t just introduce **new risks**. It multiplies the **surface area** of existing ones.

This is validated by **external signals** as well:

- **EU AI Act** readiness accelerating
- Increased board-level concern around **AI accountability**
- Regulatory pressure converging with **buyer trust expectations**

Business Continuity Moves to the Forefront

We observed a sharp increase in standalone business continuity questionnaires in 2025.

Specifically, buyers want clarity on:

- Incident response readiness
- Operational resilience under disruption
- Vendor survivability during outages, cyber events, or geopolitical instability

What it signals for **2026**:

Business continuity is no longer buried inside security reviews. It’s a primary trust signal. Buyers want assurance that vendors won’t just prevent incidents but can also withstand and recover from them.

What Buyers Are Asking About Most in 2025

Most common emerging question themes

- AI governance & model oversight
- Business continuity & resilience
- Data breaches and incident history
- MCPs and human oversight
- Customer data usage for AI training

What didn't change

There was no significant rise in new evidence formats (screenshots, attestations, reports). The transformation isn't about documentation volume. It's about depth, precision, and specificity.

Frameworks that matter (and the Ones Gaining Ground)

Most referenced frameworks (unchanged leaders):

- SOC 2
- ISO 27001
- GDPR

Fastest-growing mentions:

- EU AI Act
- DORA

Regulatory readiness and the ability to prove it quickly is a competitive advantage.

ESG Questions: A Notable Pullback

Less than 5% of questionnaires included ESG or sustainability questions.

What this suggests:

- **Economic pressure** is refocusing buyers on core operational risk
- ESG is increasingly handled **outside security workflows**
- Trust conversations are **prioritizing resilience, continuity, and AI accountability**

Scale, Efficiency, and the Reality of Modern Assurance

Questionnaire Volume Is Exploding

Security questionnaires aren't just increasing in number. They're increasing in complexity, urgency, and business impact.

Across industries, organizations are seeing:

- Year-over-year growth in inbound questionnaires
- Nearly 100 questions per questionnaire on average
- Increasing demand for expedited requests
- Clear seasonal spikes aligned with enterprise budget cycles and QBR timelines

Busiest Months

August

September

October

Late summer and early fall continue to be peak periods for questionnaire volume, as enterprise buyers finalize vendor selections and push deals across the line before year-end planning cycles. The surge is predictable but that doesn't make it easier.

What this tells us:

- Questionnaire complexity is rising
- Burnout risk for security and GRC teams is real
- Manual processes are reaching their ceiling

Trust Centers and the Maturation of Buyer Trust

Trust Centers Are Now Standard

By 2025, Trust Centers were **no longer a differentiator.** They became baseline expectation.

Nearly **99% of SecurityPal customers now maintain an internal knowledge library**, reflecting a broader shift in how buyers consume trust information. Rather than initiating lengthy email threads or one-off document requests, buyers increasingly expect trust to be **self-serve, centralized, and always up to date.**

Trust Centers have become the first line of assurance, deflecting inbound questionnaires, accelerating evaluations, and allowing security teams to focus on higher-risk, higher-impact reviews.

What buyers access most

Across SecurityPal-powered Trust Centers, the most frequently requested artifacts remained consistent:

1. SOC 2 Type II Report
2. Penetration Test Summary
3. ISO 27001 Certificate
4. SOC 2 Bridge Letter
5. AI Disclosure

The emergence of AI disclosures alongside traditional audit artifacts is especially telling. In 2025, buyers no longer viewed AI transparency as a **"nice to have."** It became a baseline expectation.

What this suggests:

AI disclosures are now as essential to trust as audit reports. Buyers want immediate clarity on how AI is used, governed, and constrained without friction or follow-up.

Agentic AI in Practice: What the Usage Data Reveals

SecurityPal Copilot usage data reveals a decisive shift: AI in assurance has moved beyond trial phases and into daily operational reliance. Copilot queries more than doubled from 2024 to 2025 across customers and internal SecurityPal teams alike.

This level of expansion signals something larger than experimentation. It reflects workflow transformation.

What this suggests:

Assurance teams can no longer keep pace with demand using manual workflows alone. AI copilots are becoming core infrastructure that support faster responses, consistency, and scalability.

This is no longer about testing AI. It's about operational dependence. The teams that can't augment human expertise with AI assistance will fall behind on speed, consistency, and trust.

What This All Means for the **Future of Assurance**

The data from 2025 points to a critical reframe: assurance is no longer a bottleneck. It's a growth function.

Security reviews now influence:

- Deal velocity and win rates
- Buyer confidence during evaluations
- Expansion and renewal conversations

As trust becomes a prerequisite for doing business, assurance teams increasingly sit at the intersection of security, sales, legal, and product.

But expectations have outpaced traditional operating models.

Why legacy approaches are failing

2025 made a several constraints unmistakeably clear:

- Manual questionnaires don't scale with **AI-driven** software growth
- Static documentation can't satisfy **continuous AI scrutiny**
- One-time reviews **fail in constantly changing** regulatory environment

Trust must now be provable, continuous, and fast, without overwhelming already-stretched teams.

Operationalizing Assurance Management at Scale

A new operating model is emerging that's designed for volume, velocity, and complexity.

Modern assurance management requires:

- **Centralized**, reusable knowledge
- **Agentic AI** to handle repeatable, high-volume work
- **Human-in-the-loop** expertise for judgment, nuance, and accountability
- **Real-time trust** signals that reflect how products actually operate

This model doesn't replace human expertise. It protects the expertise so that skilled professionals focus where they add the most value.

The Trust Stack of the Future

Across thousands of reviews, SecurityPal's data tells a consistent story:

- Trust expectations are rising
- Risk surfaces are expanding through AI-driven software creation
- Security and GRC teams are stretched thinner than ever

The path forward isn't more headcount. **It's better systems.**

By combining proprietary assurance data, agentic AI, and deep human expertise, SecurityPal is defining the future of assurance management, one that scales with innovation rather than slowing it down.

SecurityPal is the first and only platform of its kind to deliver assurance through a true AI agent + human expert model, enabling teams to move at machine speed without losing accuracy, compliance, or trust.

Trust must keep pace with change. SecurityPal makes that possible.

Trusted by the **Leading Enterprises**

OpenAI	Figma	Airtable	PLAID
orum	grammarly	stackoverflow	crunchbase
Snap Inc.	CONTENTSTACK™	MongoDB.	DOMO

Ready to Elevate Your Security Strategy?

Connect with SecurityPal expert today.

✉ contact@securitypalhq.com   Salesforce Tower, San Francisco, CA 94105   ☎ +1 650-842-0352