Artificial Intelligence You Can Trust

An Evaluation Framework for Navigating Business Stakeholders

EXECUTIVE SUMMARY

In financial services, while the promise of AI-driven transformation in risk screening is compelling, we understand that it's about control, auditability, and absolute trust.

The reality is, most "AI-powered" solutions in the market today fall short of enterprise-grade rigor. They often lack alignment with stringent model risk management protocols, struggle to articulate decision-making logic, and, crucially, fail to offer a repeatable framework for evaluation, deployment, and governance within a regulated institution.

At Sigma360, we've directly addressed this gap. We've developed a comprehensive framework specifically engineered to equip compliance leaders with a structured methodology for assessing Generative AI solutions. Whether your path involves our technology or another, our framework, called GRACE, provides the blueprint to meet regulatory and internal audit expectations, enabling risk decisions that are **fast, explainable, and inherently defensible.**

This brief outlines how GRACE empowers you to cut through the noise and integrate AI into your compliance program, without compromising control.

Risk decisions must be fast, explainable, and inherently defensible.

Why Compliance Teams Need a Framework, Not a Feature List

The market is saturated with AI vendors touting revolutionary claims for compliance. They showcase sleek interfaces, leverage buzzwords like "GPT-4 enabled," and promise the overnight elimination of manual reviews.

Yet, a critical truth often remains unaddressed: **they cannot truly explain their models' inner workings.** They typically fail to align with your organization's precise risk appetite. And when your internal teams — from Model Risk Management to Audit — pose the tough questions around validation, oversight, and auditability, clarity often evaporates.

As you know from prior technology integrations in risk-sensitive functions, initial excitement often dissipates when the rigorous checklists from internal stakeholders emerge. A flashy demonstration quickly loses its luster.

What is fundamentally required is not another catalog of AI features. It's a **framework** designed to provide clear, actionable answers to your pivotal questions:

- ☐ Can we unequivocally **trust the model's output**?
- ☐ Can we **explain this decision** to our board and, critically, to our regulators?
- ☐ Can we **configure the solution** to precisely reflect our institution's specific risk posture?
- ☐ Can we **document and prove,** with undeniable evidence, that it functions as intended and remains compliant?

This is precisely why we developed GRACE. It is not merely an acronym; it's a working tool, forged in collaboration with financial crime compliance teams, designed for the safe and effective deployment of GenAI in high-stakes environments. Because ultimately, the utility of AI is directly proportional to its governability.

The GRACE Framework: A Blueprint for Evaluating and Governing AI

Before committing to any AI solution within your compliance function, consider this: Is it architected to withstand rigorous scrutiny, or simply to impress in a sales presentation?

Most vendors prioritize speed of deployment; few build with inherent scrutiny in mind. The GRACE Framework addresses this by providing a disciplined lens through which to evaluate Generative AI. This is the same level of discipline you apply when assessing human analysts, policy amendments, or new vendor systems.

GRACE establishes **five core pillars** against which every model must be tested, ensuring confidence in your forward momentum.

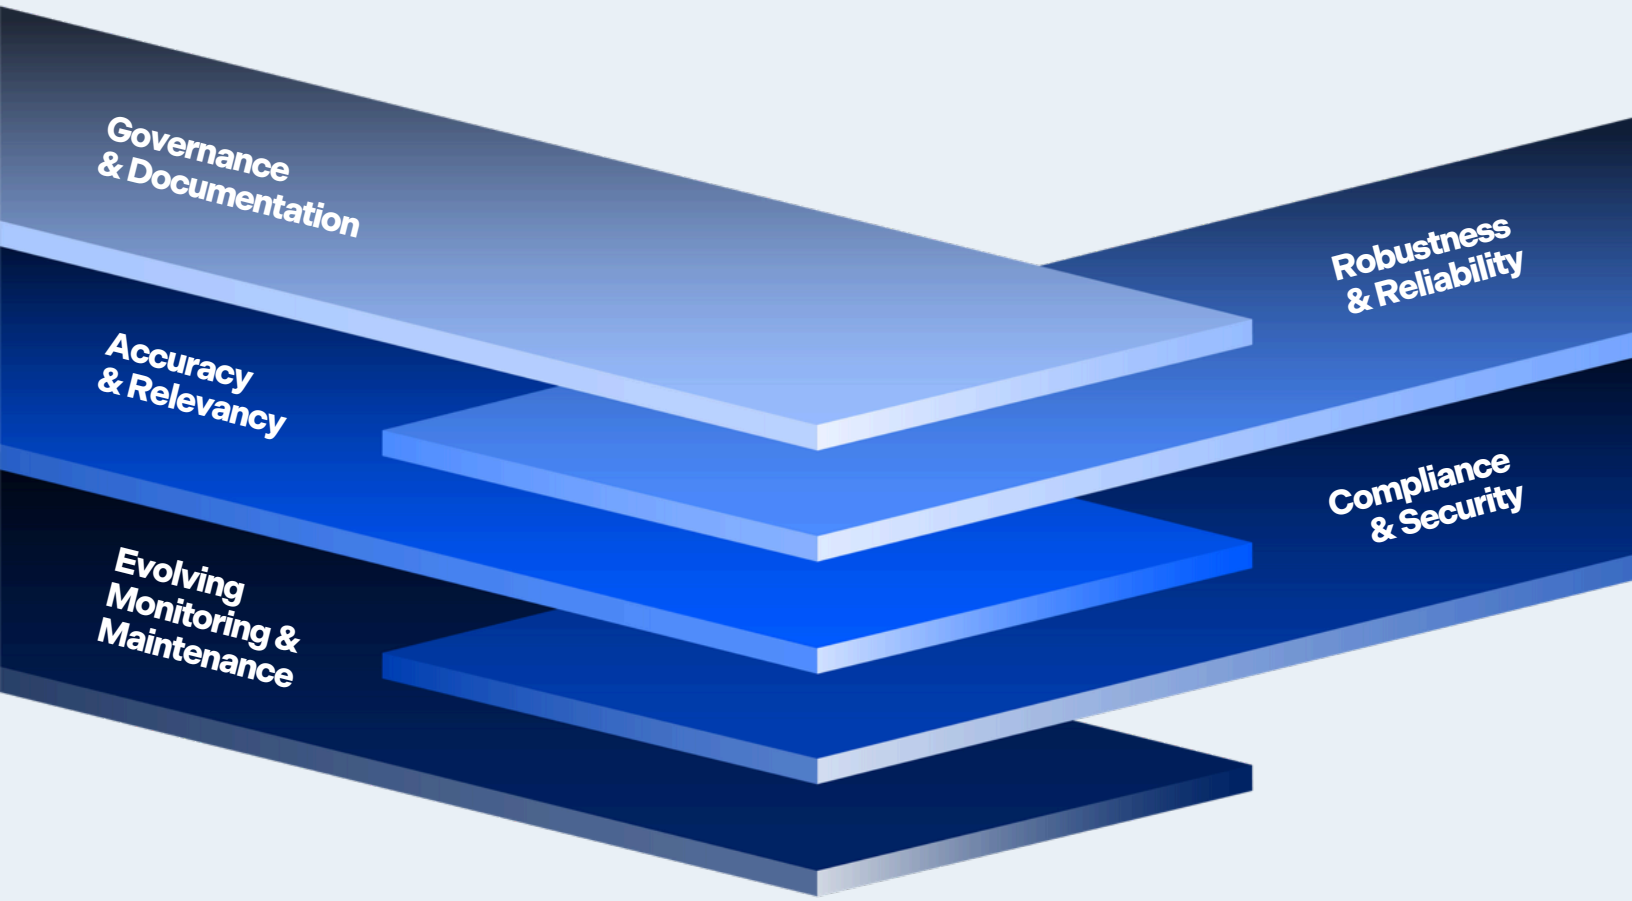

G is Governance and Documentation

If a model cannot be documented, it cannot be trusted. GRACE initiates clarity by demanding precise definition: the model's intended purpose, ownership, and the scope of its influence on decisions. Every use case mandates a clear business rationale and defined accountability. Comprehensive documentation of architecture, training data, and known limitations is required from inception. Furthermore, any model modification necessitates an explicit audit trail detailing what changed, when, and critically, *why*.

R is for Robustness and Reliability

The model's performance under duress must be unwavering. GRACE mandates **stress testing, evaluation against out-of-distribution data, and robust benchmark comparisons.** You must ascertain the model's behavior under noisy or anomalous data conditions. Equally important is its comparative performance against your existing processes. Any model exhibiting fragility under scrutiny is unfit for production deployment.

A is for Accuracy and Relevance

This pillar often distinguishes true solutions from mere technological showcases. GRACE demands **task-specific performance metrics.** Does the model genuinely differentiate between reputational and regulatory risk? Does it rigorously apply your institution's materiality thresholds? Do your expert human reviewers consistently concur with its outputs? Every model must be rigorously scored against explicitly defined business outcomes, transcending vague, generic benchmarks.

C is for Compliance and Security

Compromise is not an option here. GRACE enforces stringent controls on data privacy, usage, and storage. The model must strictly adhere to industry standards and, crucially, support **zero-data-retention policies by default.** Your teams must possess complete visibility into how the model utilizes input data, its storage location, and the robust protections safeguarding it from misuse.

E is for Evolving Monitoring and Maintenance

AI is not a static asset; its effectiveness will decay without proactive management. GRACE mandates a comprehensive plan for **continuous model performance monitoring, systematic human feedback collection, and automated triggers for necessary retraining.** This necessitates a live feedback loop, moving beyond a single quality assurance test. The objective is not merely initial efficacy, but **sustained accuracy and controlled evolution.**

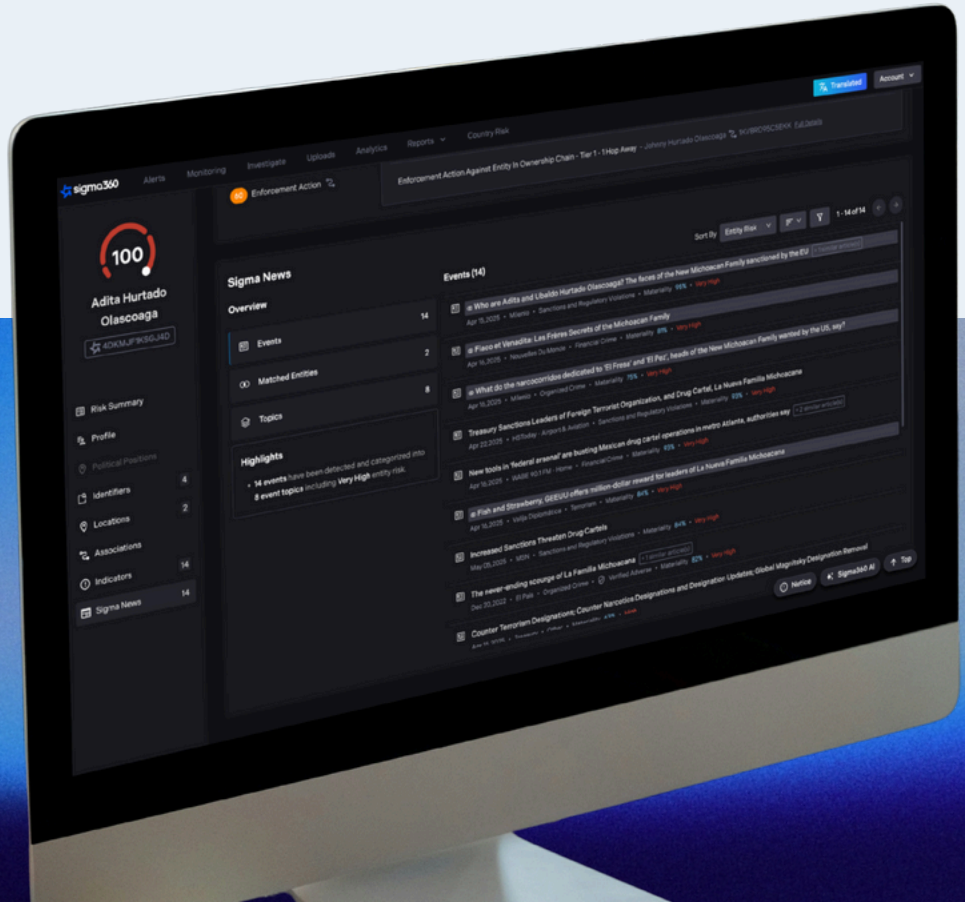

Operationalizing GRACE: A Guidance-Centric Approach

Effective frameworks transcend theoretical slides; they become actionable blueprints. GRACE is precisely that: a pragmatic guide for introducing Generative AI into your environment without compromising control, without triggering alarms from Model Risk Management, and without exposing your institution to undue regulatory scrutiny. Here's how leading compliance executives are implementing it.

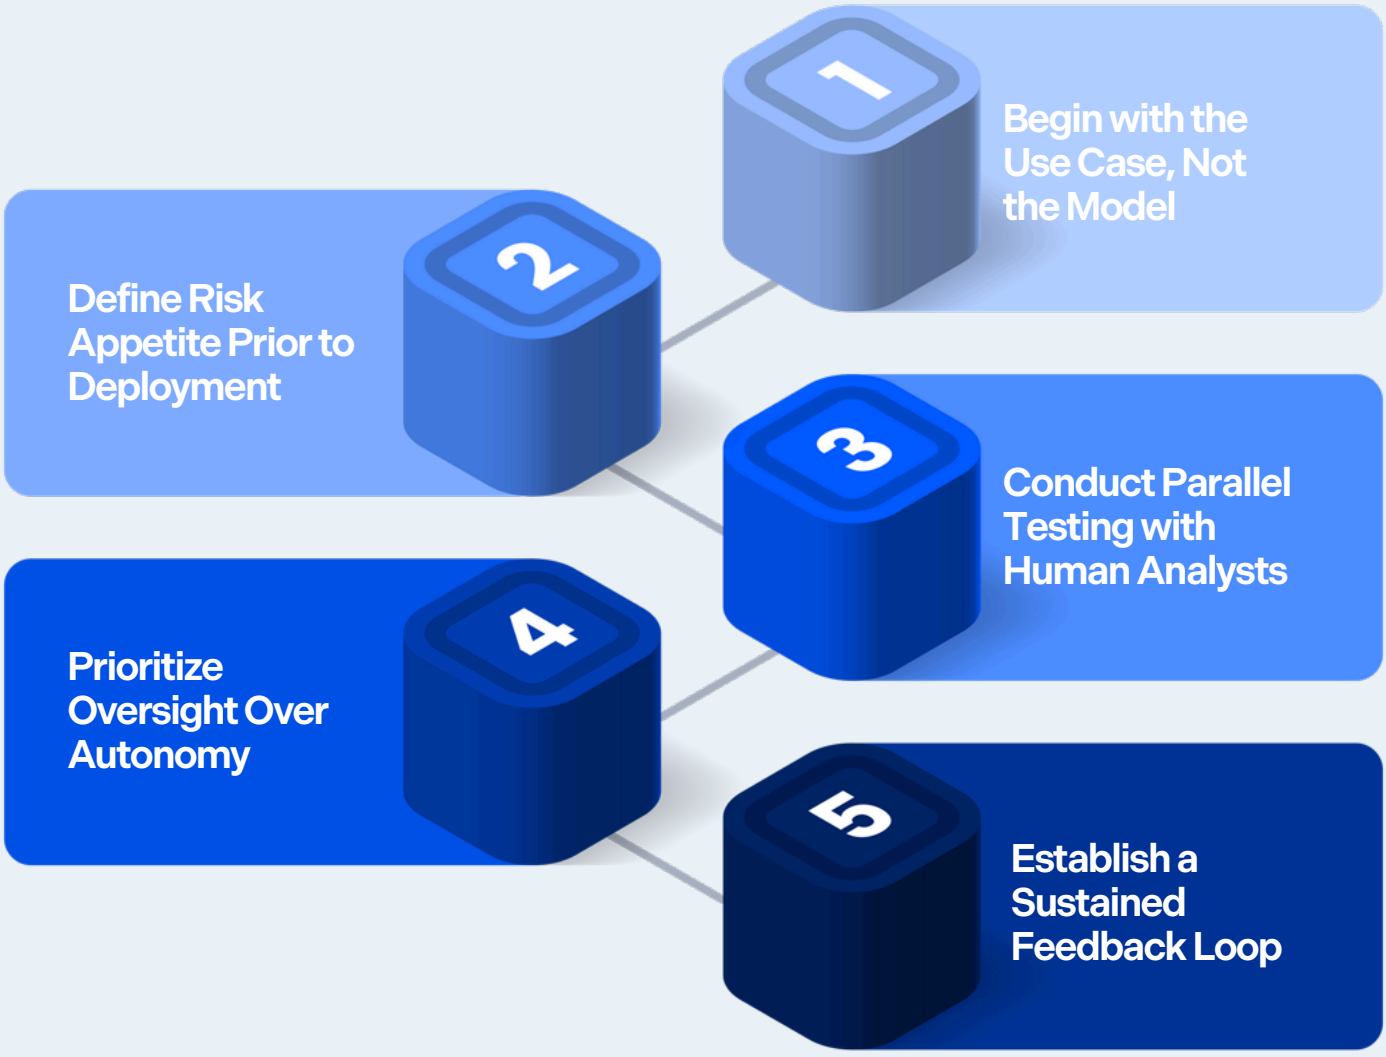

Begin with the Use Case, Not the Model

Your objective isn't merely to deploy a "chatbot." Your objective is to reduce alert backlog, generate consistent summaries, or filter out irrelevant media noise. GRACE commences with the desired outcome. What specific decision is being supported? Who is the decision-maker? What quantifiable improvement are we seeking? These fundamental questions define the AI's imperative; all subsequent steps flow from this clarity.

Define Risk Appetite Prior to Deployment

A common pitfall in pilot programs is the absence of pre-defined risk thresholds. An over-flagging model wastes analyst resources; an under-flagging model exposes the institution to critical missed risks. GRACE mandates the upfront establishment of these thresholds. What constitutes a potential match? What defines materiality? Which alert types are eligible for AI-driven auto-clearance, and which require human intervention? These are policy decisions that must be set by your leadership team, guiding the model, not being dictated by it.

Conduct Parallel Testing with Human Analysts

Before any model achieves live status, it must be rigorously tested against your existing human capabilities. Implement parallel review workflows. Quantify agreement rates. Analyze discrepancies and understand the underlying rationale. This isn't just quality assurance; it's the bedrock for building trust both within your team and with your auditors. GRACE provides specific testing protocols and prompts to ensure these comparisons yield meaningful insights.

Prioritize Oversight Over Autonomy

The ultimate goal is not automation for its own sake, but **reliable decision support**. Under GRACE, every AI action must be explainable, reversible, and demonstrably governed. If a model closes an alert, your administrator must understand the precise rationale. If it escalates, your analyst must grasp the triggering criteria. And crucially, if an error occurs, your quality assurance team must possess the mechanisms to promptly identify and correct it.

Establish a Sustained Feedback Loop

Upon deployment, the AI model becomes an integral part of your compliance team. Like any high-performing team member, it requires continuous coaching and refinement. Implement systematic triggers for retraining based on detected drift, identified errors, or evolving policy requirements. Leverage structured human feedback to incrementally enhance accuracy and relevance over time. GRACE codifies this continuous feedback as a mandatory requirement, not an optional afterthought.

AI isn't plug-and-play — it's policy-aligned, rigor-tested, and continuously supervised. GRACE turns that complexity into clarity.

A Vendor Evaluation Checklist for Compliance Teams

Deploying AI in compliance is no longer a "nice-to-have"; it is a strategic competitive advantage, contingent upon responsible implementation.

Before engaging any AI vendor, you must ascertain their capacity to meet the rigorous standards expected by your internal stakeholders: Model Risk, Audit, Legal, Technology, and critically, your own operational teams. GRACE provides the necessary structure; this checklist translates it into actionable steps.

Use the questions in the checklist on the next page to identify partners with real capabilities, not just solution sellers.

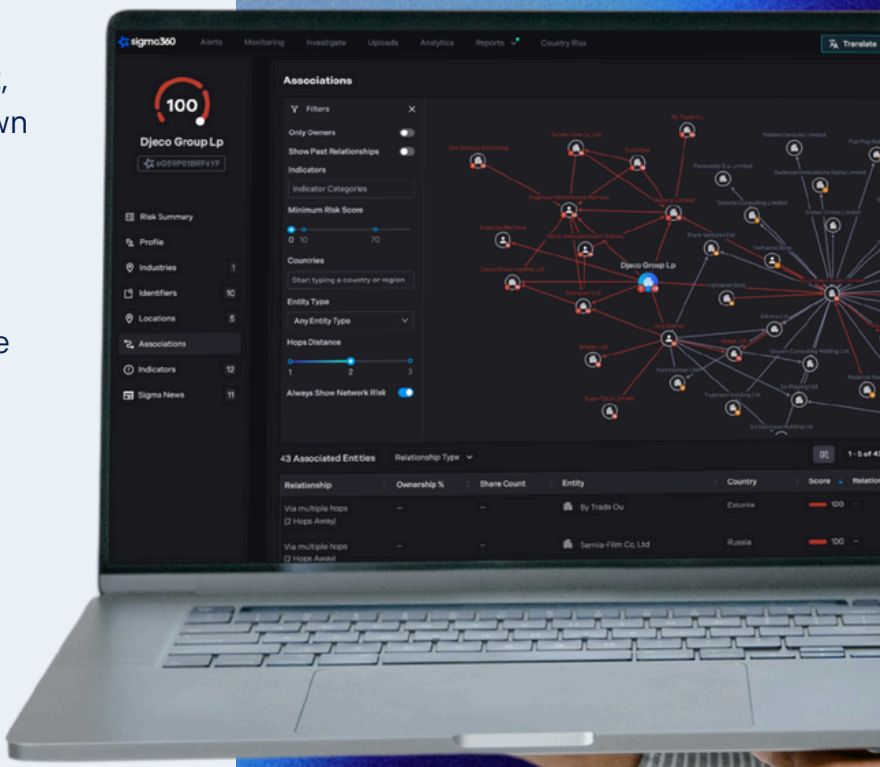

Vendor Evaluation Checklist for Compliance Teams

AI in compliance is now a strategic advantage, requiring alignment with internal standards across risk, legal, tech, and operations. GRACE provides the framework, this checklist helps identify real partners, not just solution sellers.

Governance and Documentation

- [ ] Can the vendor articulate the precise intended use case and business rationale?
- [ ] Is the model's architecture, training data, and performance rationale comprehensively documented?
- [ ] Are clear version controls and immutable audit logs provided for model changes?
- [ ] Is accountability for the model's performance and oversight explicitly assigned?

Robust and Reliability

- [ ] Has the model undergone rigorous testing under edge cases or stress conditions, mimicking real-world volatility?
- [ ] Can the vendor furnish validated performance benchmarks and comprehensive validation results?
- [ ] How does the model effectively handle noise, data duplication, or out-of-distribution data inputs?
- [ ] What are the established protocols for model failure detection and graceful degradation?

Accuracy and Relevance

- [ ] Are model outputs directly aligned to specific business goals, such as materiality thresholds or reputational risk definitions?
- [ ] Does the vendor support a robust human-in-the-loop (HITL) validation mechanism for continuous improvement?
- [ ] Are model outputs quantitatively scored, clearly explainable, and adaptable to feedback?
- [ ] What empirical evidence demonstrates quantifiable improvements in workflow efficiency or decision quality?

Vendor Evaluation Checklist for Compliance Teams

Compliance and Security

- [ ] Does the vendor unequivocally support zero data retention policies for customer input data?
- [ ] Is the solution certified and demonstrably aligned with critical regulatory standards like GDPR, CCPA, or relevant financial services mandates?
- [ ] Can the system be deployed within your preferred infrastructure model (e.g., cloud tenancy, on-premise, hybrid) with full control?
- [ ] Are data privacy, auditability, and control fundamental architectural components, or merely bolted-on features?

Evolving Monitoring and Maintenance

- [ ] What robust mechanisms are in place for continuous model performance monitoring over time?
- [ ] Are there structured feedback loops designed to systematically refine or retrain the model based on operational insights?
- [ ] Can the vendor transparently demonstrate how model updates are rigorously validated prior to release?
- [ ] What specific controls are deployed to proactively manage and mitigate AI drift?

This checklist is not a perfunctory exercise; it is your critical protection. Leverage it with Sigma360, or any prospective partner. By upholding this exacting standard, you safeguard your compliance program, empower your personnel, and maintain the agility to innovate without regret.

Moving Forward with Confidence

Today's compliance leaders bear the dual imperative of embracing innovation while rigorously defending every decision. This is an escalating challenge. Most AI vendors will readily articulate their models' capabilities. Fewer will transparently detail their governance. And fewer still will empower you to evaluate their technology against your own exacting internal standards.

GRACE is more than a framework. It represents a **commitment to architecting AI that earns its place within regulated environments.** It underpins how Sigma360 develops, tests, and validates every feature we offer. And it is how we empower compliance teams to move with decisive speed, without ever compromising on trust.

Whether you are in the nascent stages of exploring Generative AI or actively addressing stakeholder inquiries, GRACE provides the structured foundation to lead. On your terms. At your pace.

To explore a deeper dive into real-world workflows and oversight strategies tailored to your precise risk appetite, we invite you to schedule a consultation or demonstration with our team. We are prepared to help you move forward with confidence.

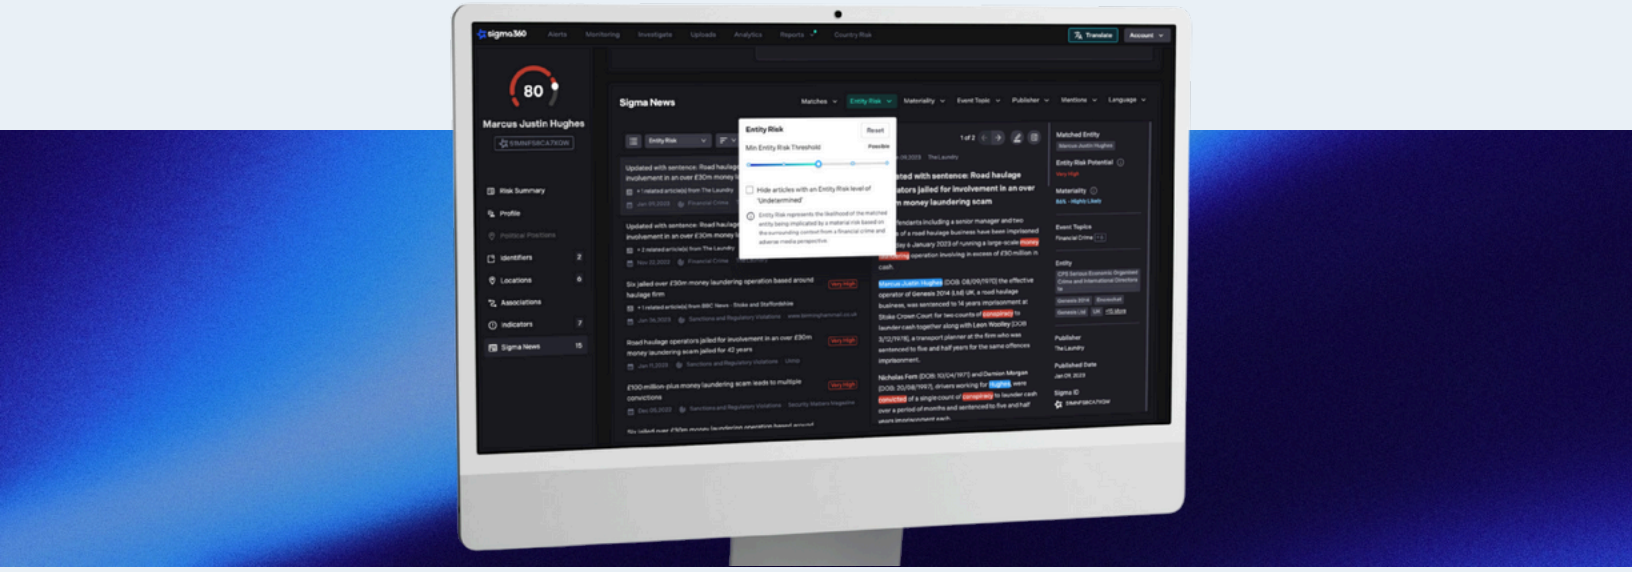

The *Future* of Financial Crime Compliance Starts with Sigma360

Sigma360 delivers a cohesive, data-centric approach to risk intelligence. Our solutions help firms identify both financial and non-financial risks in real-time at scale.

Compliance teams gain instant access to thousans of integrated data sources, and custom risk scoring, for 1 billion companies and people, makiing Sigma360 the world's most complete risk analysis platform.

Get Sigma360's AI-curated risk intelligence and taxonomies to boost your team output.

90%

Fewer Manual Match Reviews

Our AI Agent clears low-risk matches instantly, surfacing only high-risk entities for review

3x

More Data Than Legacy Providers

Covers 600,000+ publications in over 50 different languages, delivering broader, deeper visibility

95%

Less Time Reviewing Adverse Media

Adverse Media Summary filters noise and highlights only what matters, so analysts can act faster.

Sigma360 Solutions

Sanctions & Watchlist Screening
Real-time screening for PEPs, Sanctions, and Adverse Media

Adverse Media Screening
AI-enabled, continous global news screening of entities & people

Perpetual Know Your Customer (KYC)
Ongoing risk mitigation with customizable scoring and alerts

AML Investigations
Global risk intelligence centralized scoring and alerts

Enhanced Due Diligence
Global risk intelligence centralized scoring and alerts

Country Risk Ratings
Expert-designed risk-factor based view into over 260 countries

Counterparty Credit Risk
Comprehensive analysis frameworks, risk intelligence and financial data

White-Glove Onboarding and Optimizations
Receive dedicated, high-touch support from onboarding to ongoing optimizations

Need Customization?
Our team works with you to tailor our flexible features to address your unique business challenges