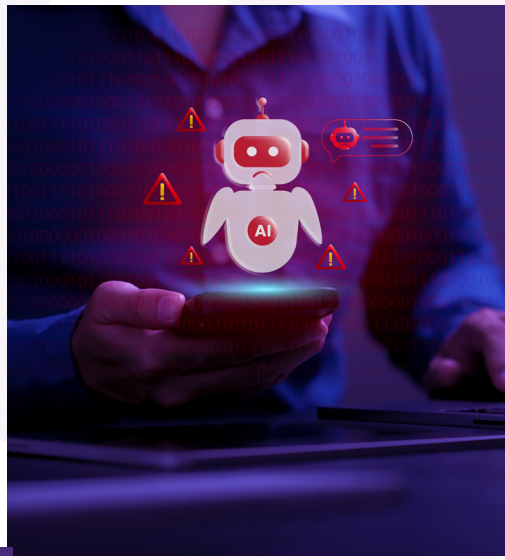
October 2026



Getting Ahead of AI Risk

5 Practical Steps for Enterprise Risk Management (ERM) Leaders

Enterprise AI adoption has outpaced most risk frameworks. Here's how to catch up.

Why This Moment

AI adoption in the enterprise has reached a tipping point. Many of your enterprise tools have AI embedded, whether you asked for it or not. Employees are using AI tools that IT hasn't approved. And the risks, including operational, regulatory, and reputational, are accumulating faster than most organizations' risk frameworks were designed to handle.

By the Numbers

66%
of employees have used unauthorized AI tools at work[1]

362
AI incidents recorded in 2025, up 55% from the year before[3]

Only 13%
of consumers completely trust AI[5]

48%
of boards that have approved major AI investments have not set AI governance expectations[4]

ERM leaders weren't necessarily asking to be thrust into the middle of this. But the speed of AI adoption, the significance of the risks, and the general lack of clear ownership and governance within most organizations means that ERM is often the best-positioned function to bring order to the chaos. No other function has the cross-domain visibility, the governance relationships, and the mandate to connect risk to strategy that ERM does.

What's required isn't a complete overhaul of how you manage risk. It's adapting your approach, recognizing that AI risk has some unique characteristics that call for a fresh look at how you identify, categorize, govern, respond to, and monitor it.

Here are five practical steps to get you started.

Step 1	Step 2	Step 3	Step 4	Step 5
Establish Your Organization's AI Risk Profile	Update Your Risk Taxonomy	Build Governance That Actually Governs	Develop Response Strategies for AI Risks	Monitor Continuously — Not Just Faster

Step 1:

Establish Your Organization's AI Risk Profile

Before you can manage AI risk, you need to understand what kind of AI risk you actually have. This sounds obvious, but it's a step that is often skipped.

The first question to answer is: what is your organization's AI posture? Are you primarily an AI user —deploying third-party tools like Microsoft Copilot or vendor AI-powered analytics? Are you an AI developer — building proprietary models or fine-tuning foundational ones? Or, increasingly likely, are you both?

The answer shapes your risk framework in meaningful ways. If you're primarily a user, your focus should center on vendor governance, model transparency, and data handling. You need to understand what the AI systems you're relying on are actually doing, and what happens when they fail. If you're a developer, the risk surface is broader: model bias, training data integrity, deployment controls, and the direct liability that comes with putting your own AI in front of customers or employees. And if you're both, you need a framework that speaks to each.

Understanding the full scope of AI in use is harder than it sounds. Research suggests that two-thirds (66%) of professionals have used unauthorized AI tools at work[1], and 45% of workers find workarounds when their company blocks an AI application[2]. But the line between shadow AI and sanctioned AI is increasingly blurry. Much of what gets labeled "shadow AI" is simply AI capability embedded in tools employees were already given (e.g., Copilot in Word, AI-assisted search in a CRM, recommendation engines in an ERP). The real question isn't just "who's going rogue?" It's "how much AI are we already running, how much of it do we actually understand, and do we have the right guardrails in place?"

Shadow AI

An ERM leader at a financial services firm put the problem in concrete terms during a recent Uniqus ERM roundtable conversation. The concern wasn't employees downloading unauthorized tools on the side. It was the AI already embedded in systems the organization had already approved: Salesforce, Workday, Zendesk. A vendor update, a feature toggle flipped by a local admin, and suddenly AI is processing company data inside a platform nobody thought to include in a risk assessment.

> *Do employees even know it's been turned on?*
> *- ERM leader, financial services firm.*

That question gets to the real challenge. Shadow AI today is less about rogue behavior and more about the limits of what your organization can actually see.

Completing this profile also means defining your organization's AI risk appetite — how much AI-related uncertainty you're prepared to accept in pursuit of your strategic goals. This shapes everything downstream: what risks you choose to mitigate, what you're willing to accept, and how aggressively your organization can pursue AI-driven opportunities.

The Trust Problem

Only 13% of consumers completely trust AI[5], and 82% view AI-related data loss of control as a serious threat[6]. As AI becomes more embedded in how organizations interact with customers, make decisions, and deliver services, those trust deficits translate directly into business risk. 50% of US consumers say they would prefer to give their business to brands that don't use generative AI in customer-facing content[5]. Managing AI risk well isn't just about compliance or avoiding incidents — it's about being the kind of organization that customers, employees, and partners feel comfortable trusting with AI.

13% of consumers completely trust AI

82% view AI data loss of control as a serious threat

50% of US consumers prefer brands without generative AI in customer-facing content.

Finally, define what AI risk means for your organization (including the upside). The risk of moving too slowly on AI adoption is just as real as the risk of moving too fast. A well-designed AI risk framework doesn't just protect against harm. It gives your organization the confidence to adopt AI thoughtfully, which is itself a competitive advantage.

Step 2:

Update Your Risk Taxonomy

Thinking through where AI risk fits in your overall risk landscape is critical to understanding how risks interconnect, what's driving them, and who owns them. The goal isn't to rebuild your taxonomy from scratch. Most of the risk categories you already have are the right ones. AI is reshaping them, not replacing them. Think about AI risk in two ways:

New Risks
Certain AI risks are genuinely new risk types that belong in your taxonomy.

Existing Risks AI Amplifies
Certain AI risks act as new drivers that change the nature of risks you already track.

Organized across the FACTORS framework, here is how AI risk shows up across each domain:

Domain	New AI Risks	Existing Risks AI Amplifies
Financial	• Model errors in pricing, credit, and forecasting • AI-related litigation and regulatory penalties • Valuation of AI assets	• Financial fraud via AI-powered schemes • Financial projections distorted by model overconfidence
Associates (HR)	• Bias in AI-driven hiring and performance management • Over-dependence on AI eroding critical skills	• Talent retention pressure as demand for AI-savvy talent intensifies • HR compliance exposure from automated decisions
Compliance	• AI regulatory compliance (EU AI Act, SEC disclosure, state-level legislation) • Model explainability obligations	• Privacy violations from AI data ingestion • Existing compliance gaps amplified by AI-driven scale
Technology	• Model risk (hallucination, bias, drift, opacity) • Data poisoning • AI infrastructure concentration risk	• Cybersecurity (AI-enhanced attacks, deepfakes, social engineering) • IT resilience and continuity
Operations	• Fully automated process failures with no human fallback • BCP gaps for AI-dependent processes	• Operational disruptions propagating faster and further • Quality control failures in AI-assisted work
Relationships (3rd Parties)	• Vendor-embedded AI failures • AI risk in the supply chain • Third-party model governance gaps	• Third-party due diligence gaps • Contract exposure • Vendor concentration risk
Strategic	• Disruption risk from competitors using AI • AI provider concentration • Reputational risk from AI failures • Risk of under-adopting AI where competitors do not	• ESG risk from AI deployment choices • Strategic misalignment between AI use and stated values

The harder and more important question is how to reflect AI as a driver of existing risks. Simply listing AI risks separately isn't enough, you risk creating a parallel track that leadership treats as an IT issue. Instead, consider these approaches:

01

Flag AI-driven risks on your risk register:

An asterisk or dedicated "AI Driver" column signals that these risks have a new or amplified dimension.

02

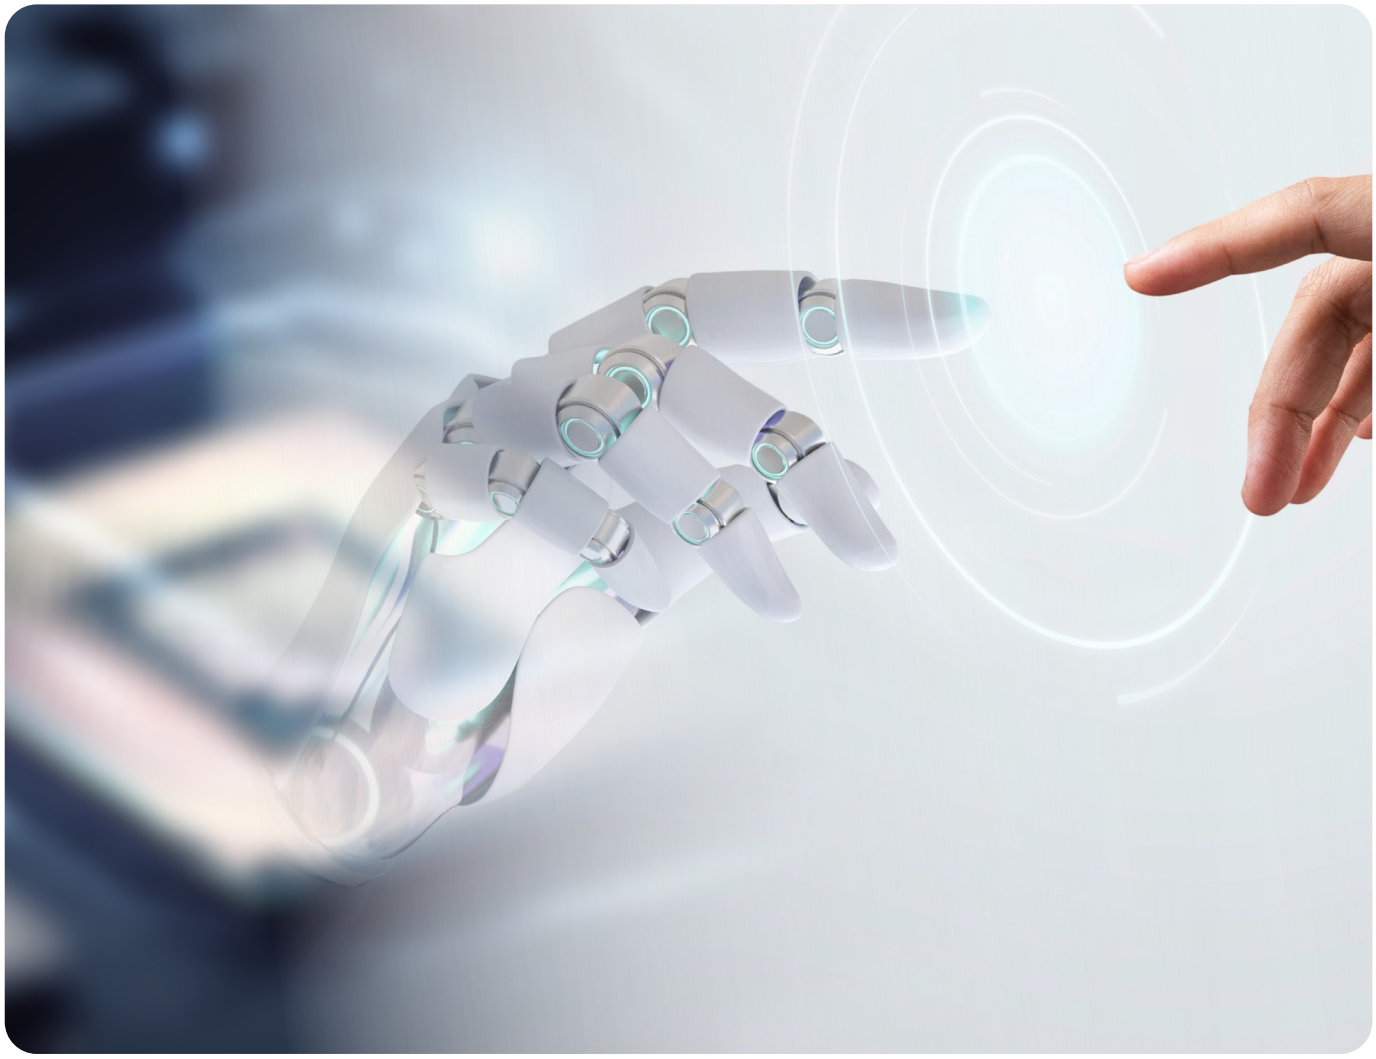
Update the drivers for each affected risk to include AI-specific sub-drivers:

For example, model drift as a driver of operational risk, or AI-generated content as a driver of reputational risk.

03

Introduce an "AI-amplified" lens in your top risk reporting:

Not a separate category, but a filter that helps leadership see which existing exposures are moving faster because of AI.



Step 3:

Build Governance That Actually Governs

Most organizations that say they have AI risk governance have a committee. What they often don't have is clarity on who can actually make a call.

The first governance question to answer isn't "who should be on the committee?" It's "who is accountable for what?" That means defining, risk by risk, who owns the exposure, who sets the appetite, and who has the authority to approve, modify, or halt an AI deployment. Without that, a committee becomes a forum for discussion rather than a mechanism for decision-making.

The second question is where AI risk governance sits structurally. There's no universally right answer, but there are wrong ones.

- Putting AI risk governance entirely under the CISO assumes AI is primarily a cyber problem — it's not.
- Embedding it within a Chief AI Officer's remit risks having the function that champions AI adoption also be the one assessing its risks.
- A cross-functional model, with the CRO or ERM function providing the coordinating layer, tends to work best.

One model worth considering is having two separate but connected bodies:

AI Risk Committee	AI Strategy Committee
Focuses on identifying and responding to the downside.	Focuses on ensuring the organization captures the upside.

Having the CRO sit on both is important. It ensures risk considerations are embedded in strategic AI decisions, and that the risk function is connected to where the real opportunities are being weighed. Without that connection, risk management becomes a veto function rather than a strategic enabler.

The Committee That Couldn't Decide

One investment management firm formed its AI Strategy Committee in early 2025 with a mandate to oversee the firm's AI use. By year-end, the committee had met four times, produced a governance charter, and had not approved or rejected a single initiative. When a portfolio team asked for clearance to use an AI tool to support financial model building, the request cycled through two committee meetings without resolution. No one escalated it. No one closed it. It waited.

> *We created the committee because the CEO expected us to," one member of the committee noted. "But we never answered the basic question: who actually has the authority to say yes?*

The committee was real. The governance was not.

For the committee itself: *resist the instinct to make it large for the sake of coverage. A smaller group of the right people (e.g., those who want to be there, who understand both risk and AI, and who have real organizational influence) will outperform a 20-person body where half the members are voluntold and disengaged. That group needs clear escalation paths to the board and a defined cadence that reflects the speed at which AI risk can develop.*

Step 4:

Develop Response Strategies for AI Risks

Identifying AI risks is necessary but not sufficient. For each material risk in your taxonomy, you need a defined response — proportionate, practical, and assigned to someone.

Think about responses across four options:

Response	What It Means	AI Example
Reduce	Implement controls to lower the likelihood or impact of the risk	• Mandatory human review of AI outputs that inform high-stakes decisions • Model validation requirements before deployment
Transfer	Shift the financial exposure to a third party	• AI-specific representations and warranties in vendor contracts • Verifying that cyber insurance policies cover AI-related incidents (many don't, yet)
Accept and monitor	Maintain current controls and monitor the risk for changes	• Documented decision to proceed with a new AI tool while acknowledging known model limitations, within defined appetite thresholds
Avoid	Decline to take on the risk entirely	• Declining to fully automate a decision process where errors carry unacceptable legal, reputational, or financial consequences

A few areas warrant additional attention beyond the table above.

Human on the loop.

Rather than revisiting the well-established principle of keeping humans in the loop, the more nuanced question today is how much of a role humans play. The conversation in most mature organizations has moved from "human in the loop" (where a human approves every AI output) to "human on the loop" (where humans supervise AI systems and can intervene), but aren't reviewing every individual decision. Getting explicit about which model applies to which process, and why, is where the practical work is. The answer will differ by use case, risk level, and regulatory context.

AI incident response.

An AI incident isn't like a typical IT outage. The failure mode may be subtle — a model producing biased outputs for months before anyone notices. The damage may have already spread widely before it's detected, and the remediation may require retraining or rolling back a model rather than just restoring a system. Your incident response protocol should define: what triggers an AI incident declaration, who has authority to pause or roll back a model, how you assess the scope of affected decisions, how you communicate with affected parties, and what post-incident review looks like. Most organizations have none of this defined.

The response to opportunity.

ERM can and should help the business define the conditions under which it's willing to move forward with AI despite uncertainty: what level of model accuracy is acceptable for a given use case, what processes are off-limits for full automation, and what governance steps are required before a new AI tool is deployed at scale. This isn't gatekeeping. It's giving the business a clear path to yes.

Step 5:
Monitor Continuously — Not Just Faster

Traditional ERM runs on a cycle. AI risk does not. A model can drift, a regulation can change, or a vendor can make a product update that materially changes your risk profile, all between your last assessment and your next one.

The answer isn't simply to run your existing processes more frequently. It's to build always-on sensing mechanisms alongside your periodic assessments:

Know what to watch.

Track model performance metrics, including accuracy, drift indicators, and bias measures, over time, and set alert thresholds that trigger a review when performance degrades. This requires coordination with technical teams, but the data should flow into risk monitoring, not just an engineering dashboard.

Build a regulatory scanning process.

The EU AI Act alone is generating significant implementing guidance and enforcement actions in real time. A manual quarterly scan won't keep up. Assign ownership for AI regulatory monitoring and define how new developments get triaged and escalated.

Define escalation triggers.

Establish the conditions under which an AI risk event automatically surfaces to senior leadership or the board, without waiting for the next reporting cycle. These triggers should be specific, not "significant AI incident", but "AI system produces output affecting more than X customer decisions" or "model accuracy drops below Y threshold."

Treat your AI risk appetite as a living document.

Established in Step 1, your risk appetite should be formally reviewed at least annually and updated out-of-cycle when something material changes (e.g., a new regulatory requirement, a high-profile AI incident in your industry, or a significant shift in your organization's AI posture). What you were willing to accept eighteen months ago may not reflect where you are today.



AI RISK DIAGNOSTIC CHECKLIST

Use this to assess the current state of your AI risk program. A “no” or “unsure” on any of the following is a gap worth addressing.

Risk Profile

Step 1

- ☐ We have mapped where AI is in use across our organization, including third-party and embedded AI.
- ☐ We have distinguished between our AI user risk and our AI developer risk (if applicable).
- ☐ We have defined our organization’s AI risk appetite.
- ☐ Our risk leadership has had an explicit conversation about the risk of under-adopting AI, not just over-adopting it.

Taxonomy

Step 2

- ☐ Our risk register includes AI-specific risk categories (e.g., model bias, data integrity, AI regulatory compliance).
- ☐ We have reviewed existing risk categories to identify where AI is a material driver or amplifier.
- ☐ AI-amplified risks are flagged or identified distinctly in our risk reporting.
- ☐ Our third-party risk framework includes questions about how vendors use AI in what they deliver to us.
- ☐ Our taxonomy generally reflects different domains that cover these FACTORS categories: Financial, Associates, Compliance, Technology, Operations, Relationships, and Strategic.

Governance

Step 3

- ☐ There is a named individual accountable for each material AI risk in our register.
- ☐ We have a defined process for approving the deployment of a new AI system or tool.
- ☐ There is a mechanism to pause or roll back an AI system if it produces harmful output, and someone with the authority to invoke it.
- ☐ We have considered whether separate AI Risk and AI Strategy committees are appropriate for our organization.
- ☐ Our board receives regular reporting on AI risk.

Response

Step 4

- ☐ We have documented response strategies (reduce, transfer, accept, avoid) for our highest-priority AI risks.
- ☐ We have defined where human-on-the-loop supervision applies and where stricter human-in-the-loop approval is required.
- ☐ Our vendor contracts include AI-specific representations or protections.
- ☐ We have an AI-specific incident response protocol.

Monitoring

Step 5

- ☐ We are tracking model performance metrics (accuracy, drift, bias) as part of ongoing risk monitoring.
- ☐ We have a structured process for monitoring AI regulatory developments on a continuous basis.
- ☐ We have defined trigger conditions for escalating AI risk events outside the normal reporting cycle.
- ☐ We review and update our AI risk appetite on at least an annual basis.

Total ‘no’ or ‘unsure’ answers:

_____ of 22.

Discuss your results with Uniqus

Share your checklist results with our team to understand where AI risk sits in your organization and what to address first.

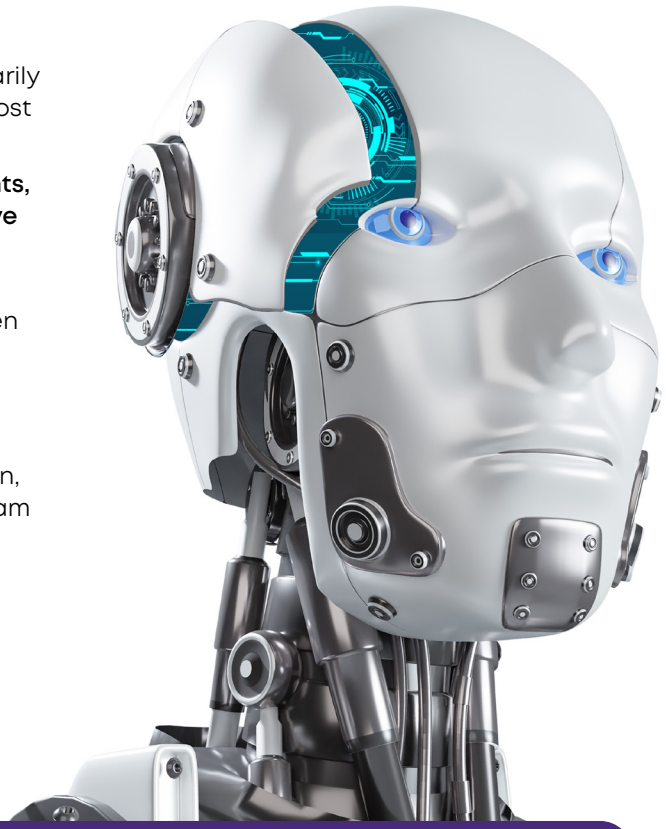
CLOSING

The organizations that will navigate AI risk best won't necessarily be the most cautious or the most aggressive. They'll be the most prepared. And there's a meaningful gap right now:

While three in four boards have approved major AI investments, Nearly half have not set AI governance expectations and have not integrated AI risk into their ongoing oversight.[4]

Reported AI incidents rose 55% in 2025 alone.[3] The gap between where organizations think they are on AI risk and where they actually are is closing, but not always in the right direction.

For ERM leaders, this moment is also an opportunity. AI risk is high-stakes, cross-functional, and genuinely difficult to govern, which is exactly the environment where a well-run ERM program proves its value. The organizations that get this right will have ERM at the center of their AI strategy: not as a speed bump, but as the function that gives the business the clarity and confidence to move forward. That's a different conversation than ERM has historically been invited into, and it's worth showing up for.



Connect with us

AI adds new risks and changes how existing ones behave. Talk to our team about what this means for your organization and where to focus first.

Matt Solomon
US ERM Leader
mattsolomon@uniqus.com

Sources

1. PagerDuty, Shadow AI Workplace Survey (2026): two-thirds of office professionals have used unauthorized AI tools
2. UpGuard, The State of Shadow AI (2025): 45% of workers find workarounds to access blocked applications
3. Stanford HAI, 2026 AI Index Report; AI Incident Database (2026): 362 AI incidents recorded in 2025, up 55% from 2024
4. Grant Thornton, 2026 AI Impact Survey (April 2026): three in four boards have approved major AI investments; 48% have not set AI governance expectations; 46% have not integrated AI risk into ongoing oversight
5. Klaviyo, AI Consumer Trends Report (2026): 13% of consumers completely trust AI; 50% prefer brands not using generative AI in customer-facing content
6. Relyance AI, Consumer AI Trust Survey (2025): 82% of consumers view AI data loss of control as a serious threat

About Uniqus Consultech:

Uniqus Consultech is an AI and global tech-enabled consulting company that specializes in Accounting & Reporting Consulting, Governance, Risk & Compliance, Sustainability & Climate Consulting, Tech Consulting, and Valuations. The Company is co-founded by consulting veterans Jamil Khatri and Sandip Khetan and backed by marquee investors such as Nexus Venture Partners, Sorin Investments, and UST.

Uniqus has a global team of 800+ professionals, led by 100+ Partners and Directors, across 13 offices in the USA, the Middle East, and India. The company serves more than 400 clients, including marquee names in each of the markets it operates in.

Uniqus is committed to leveraging technology and an integrated global delivery model to provide best-in-class consulting services to its clients.

About Our Governance, Risk & Compliance (GRC) Services:

Our Governance, Risk & Compliance (GRC) practice partners with organizations as they continue to deal with the dynamic governance, risk, and compliance landscape and navigate today's complex business environment. We seek to prepare business leaders and their teams to identify and manage risks, strengthen governance across levels, drive end-to-end compliance, design and test internal controls, anticipate and assess any disruptions, and help build business resilience. We help to deploy platforms aimed at digitization, automation, and AI-enablement of an organization's GRC functions.

A Team That You Can Trust To Deliver



Jamil Khatri
Co-Founder & CEO



Sandip Khetan
Co-Founder, Global Head of Accounting & Reporting Consulting



Nagaraj Uchil
Partner, Global Head of GRC



Matt Solomon
Managing Director, GRC

Scan this code to find more content like this

Visit us at www.uniqus.com
or follow us on