OneTrust

The OneTrust 2027 Predictions Report

Architecting AI Governance:

Blueprints for a Moving Maze

Report

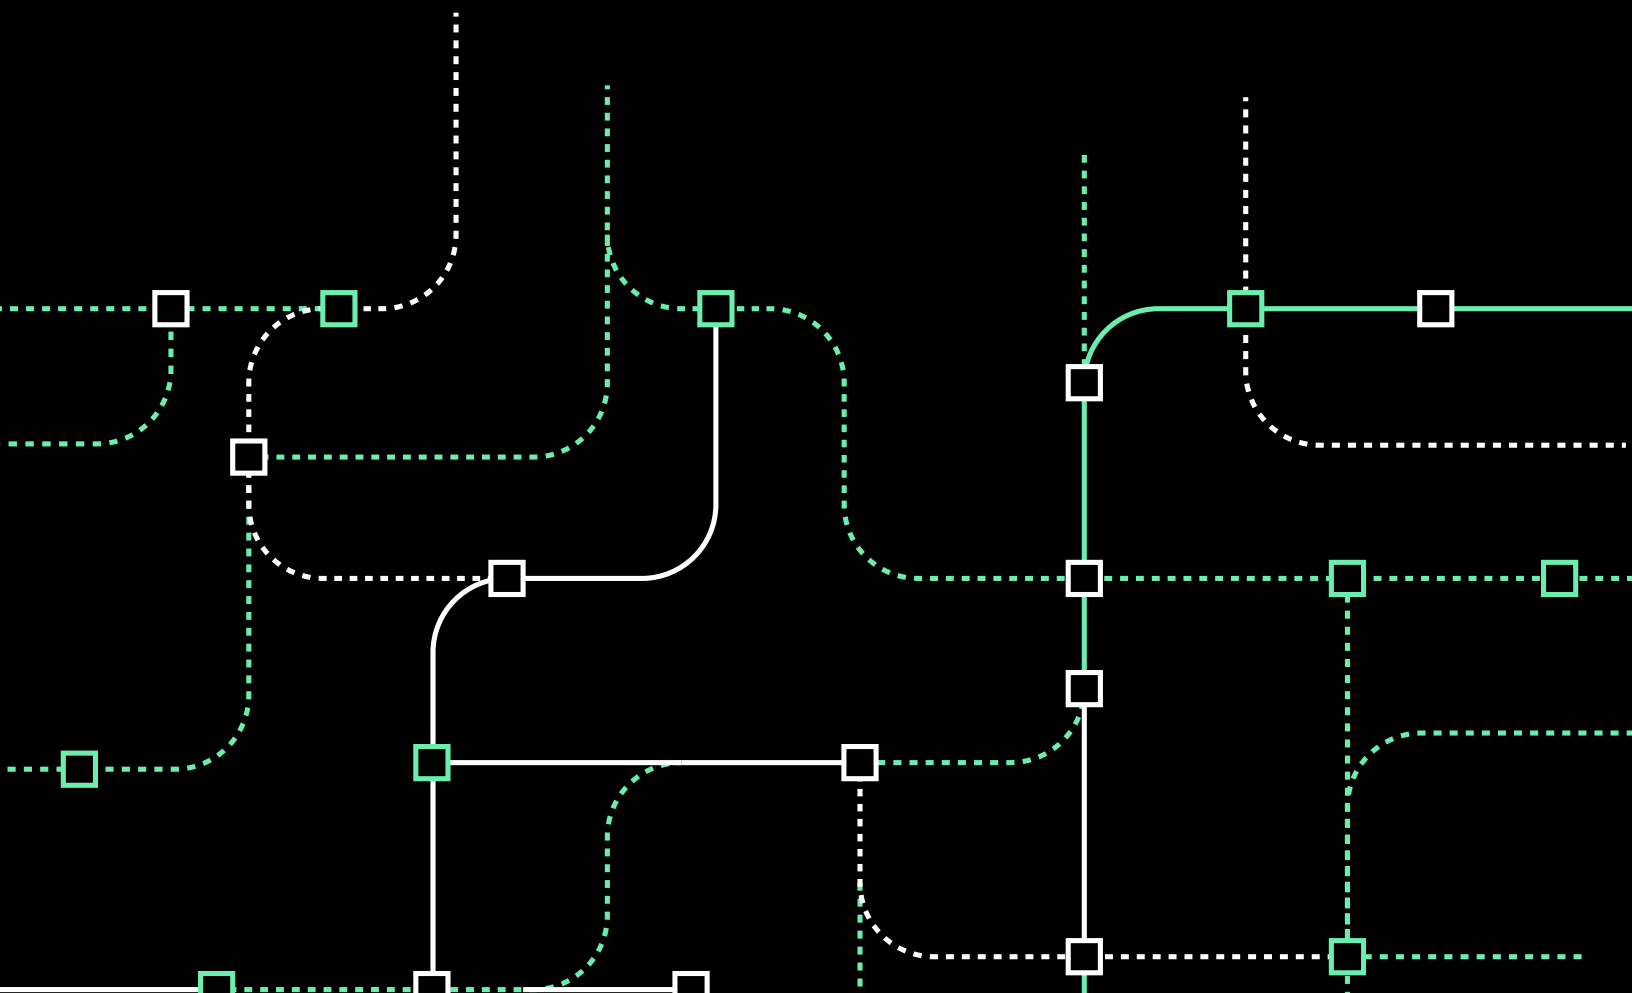


Table of Contents

Introduction: Access is granted. Authority is bounded. Trust is proven. 3

Prediction 1: No proof, no passage. AI assurance becomes a commercial gate 4

Prediction 2: Green lights can lead to the wrong destination. An approved agent can still do an unauthorized job 7

Prediction 3: The AI estate outgrows the approval. Material AI changes become new assurance events 10

Prediction 4: Permission is not an all-access pass. AI conversations make data-use choices part of the journey 14

AI regulation in 2027, as told by the experts 17

Anti-predictions: 3 things that will not happen in 2027 25

Disclaimer
No part of this document may be reproduced in any form without the written permission of the copyright owner. The contents of this document are subject to revision without notice due to continued progress in methodology, design, and manufacturing. OneTrust LLC shall have no liability for any error or damage of any kind resulting from the use of this document. OneTrust products, content and materials are for informational purposes only and not for the purpose of providing legal advice. You should contact your attorney to obtain advice with respect to any particular issue. OneTrust materials do not guarantee compliance with applicable laws and regulations.

Copyright © 2026 OneTrust LLC. All rights reserved. Proprietary & Confidential.

Introduction

Access is granted. Authority is bounded. Trust is proven.



Blake Brannon
Chief Innovation Officer
OneTrust

The debate about the pace of frontier AI raises an immediate enterprise decision: what are we prepared to let these systems do? That question becomes consequential when an agent can change a business record, use customer information, or initiate work in another system.

We already know how to answer it for people. A new hire gets a badge, is told what they may decide, and earns trust over time. Agents need the same three things, and each one works differently than it does for a person.

An agent gets its access the moment someone connects it to a system, and often it inherits whatever the person who built it could reach, which is frequently more than the task needs. That part is easy and it is already happening. In OneTrust's 2026 AI-Ready Governance Survey, 87% of respondents say their organizations encourage agent use, including 47% with clear governance and controls and 40% whose capabilities are still developing. Encouraging use is the right instinct. The exposure is in what gets assumed along with it. Access is granted.

What an agent may do with that access is a separate decision, and someone has to make it. Access is what a system can reach. Authority is what it may do with that reach, for what purpose, and up to what limit. An agent with access to the order system may be able to retrieve order records. Its authority might allow refunds up to $200 on orders less than 30 days old, with anything beyond those limits referred to a person. Those limits have to be written down and enforced where the refund actually happens, not implied by the login, because an agent's job widens quietly. It picks up another connection, learns something in a conversation that a form would never have asked, and puts it to use beyond the original purpose, all under the same approval. Nothing in the agent will flag that on its own. Only 5% of respondents report clear coordination and accountability across the AI lifecycle, and without it, individually permitted steps can combine into an outcome nobody authorized. Authority must be bounded.

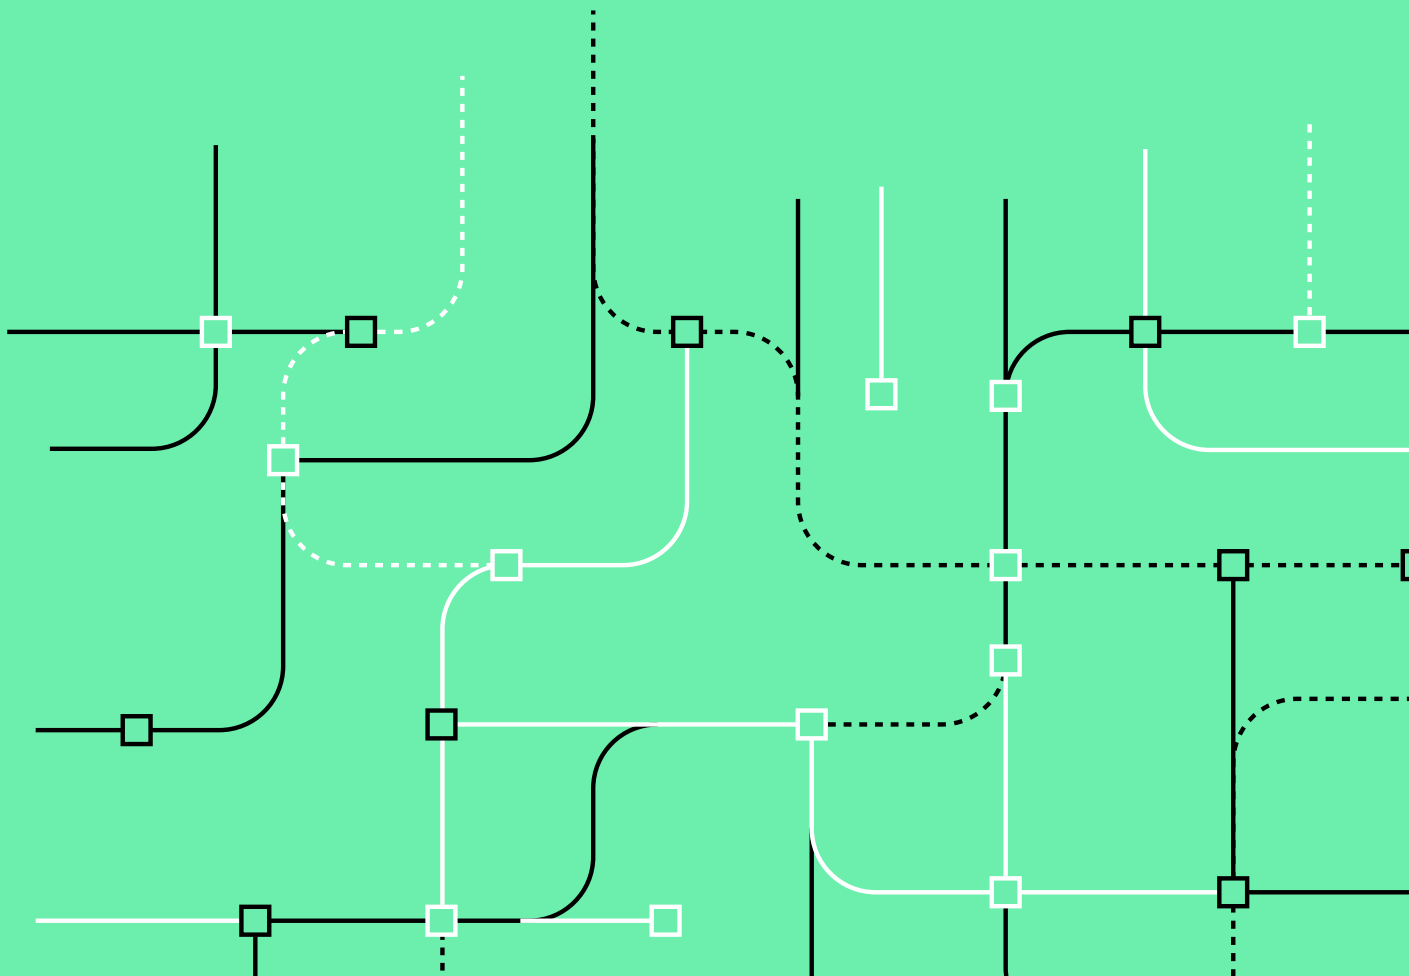
For an agent to be trusted, it has to show over time that it did what it was allowed to do and nothing else, and someone other than the agent has to be able to check. An agent can run under a person's credentials and carry their name through every system it touches, but it does not carry the track record that made those credentials trustworthy. Identity is inherited. Trust is not. And the agent cannot be its own witness; its account of what it did is one more output from the system under review. The record has to be kept outside the agent's control and retained as evidence that can be reviewed independently. Only 28% of respondents report governance evidence and audit trails, the least common activity in the survey. The first prediction in this report shows who will ask for that evidence, and when. Trust must be proven.

None of this argues for slowing down. Clear limits remove avoidable delay, and governance exists to make the harder calls at the right speed and to name who owns the risk that remains.

The four predictions that follow examine how access, authority, and evidence will shape adoption and business outcomes in 2027, from commercial assurance to the limits of delegated work, the changing AI estate, and the customer's own data. Before expanding an agent's authority, can your organization explain its permitted purpose, stop an unauthorized action, and reconstruct what happened? That is the test this report asks leaders to prepare for.

Prediction 1

No proof, no passage. AI assurance becomes a commercial gate



Prediction 1

No proof, no passage. AI assurance becomes a commercial gate

In 2027, buyers of consequential AI-enabled services will increasingly require current, deployment-relevant evidence before activating or expanding their use. The assurance required will reflect the capability's autonomy, business impact, and potential harm. That pressure may extend into contracting, renewal, and budget decisions.

Consider a supplier releasing a new AI feature. Development is complete, the capability is available, and the customer can turn it on. Yet the customer waits. The unresolved question is whether the assurance already provided covers what the service can now do.

For the supplier, technical readiness has not yet translated into customer acceptance. The commercial value of the release depends partly on resolving the risk questions that stand between availability and activation.

We expect this connection to become more explicit during 2027. Buyers will ask whether existing assurance applies to the capability, data connections, and deployment they're being asked to accept. Missing or outdated evidence may delay activation, narrow the scope of use, or restrict expansion.

OneTrust research shows an uneven evidence foundation: 28% report governance evidence and audit trails, the least common activity listed. Meanwhile, 38% prioritize preparing for regulations, standards, or audits over the next 12 months.

> "Customers need assurance that reflects the capability they are being asked to use. When functionality changes, earlier testing may leave new questions unanswered. Relevant evidence helps them understand the exposure and decide whether they are comfortable enabling it.",

Tim Mullen
Chief Information Officer
OneTrust



Evidence must help resolve a decision. A penetration test identifies vulnerabilities within its scope. A control evaluation tests whether a delegated limit holds. Buyers and business owners still need to know what was examined, which protections remain effective, and who can accept the remaining risk.

Recent Gartner research identifies AI governance as an emerging buying criterion, describing customer requests for evidence of traceability, monitoring, and override mechanisms. The analysis also connects incomplete answers with longer due diligence and potential insurance constraints. This supports the direction of the forecast, although it does not establish how widely those demands already determine commercial outcomes.

The pressure also reaches executive teams. Gartner argues that leaders need to understand which business decisions AI influences and who is accountable for the consequences. We expect that scrutiny to make evidence more consequential when organizations decide whether to activate a capability or expand its authority.

AI-assisted supplier screening is a further possibility for 2027. Published policies and assurance materials could influence early screening. Whether that changes shortlisting decisions at scale remains an open question for 2027.

For Chief Financial Officers and investment owners, the related question is whether a capability delivers enough value to justify its operating cost and exposure. We expect decisions about further investment to draw more explicitly on those measures alongside assurance. A capability may remain within its approved limits while still failing to justify expansion.

For suppliers, resolving assurance questions becomes part of getting an AI capability into use. Commercial teams need to identify which evidence is missing and who can resolve the decision. For buyers, the practical question is whether to activate the capability, impose conditions, or wait. [

What could change the forecast

Buyers may continue accepting familiar questionnaires and periodic testing. Insurance exclusions can shift exposure without prompting governance investment, while litigation develops unevenly. The forecast weakens if fresh evidence rarely changes activation, terms, or expansion.

Probability: Moderate.



Industry impact:

High where assurance affects acceptance or expansion. The consequence can reach product adoption, commercial terms, and the confidence to commit resources.

Priorities for governance leaders

- Security and product: Match assurance to the capability and version being offered.
- Commercial teams: Track requests that delay acceptance or change deployment scope.
- Risk and legal: Identify relevant contractual, coverage, and disclosure questions.
- Finance and business owners: Define the expected value, operating-cost, and assurance evidence required before further investment or expansion.

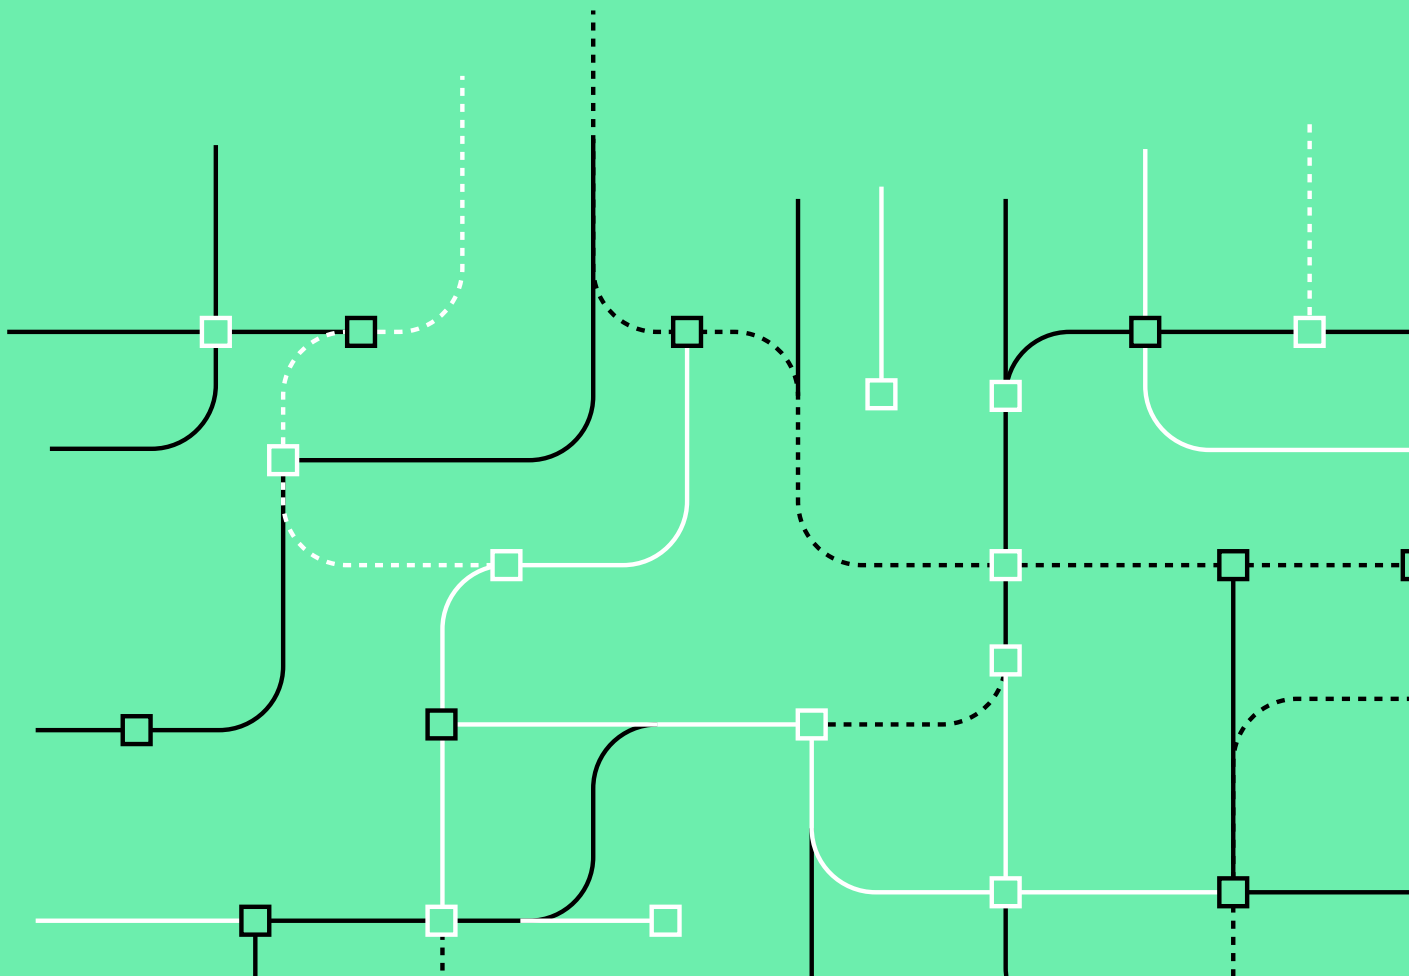
What to look for in 2027

- Buyers withholding activation or expansion until fresh evidence is provided.
- Contract provisions requiring assurance after material AI changes.
- AI-specific underwriting questions connected to coverage terms or exclusions.
- Funding and expansion decisions that explicitly consider value delivered and operating cost alongside control evidence.

[1] *When AI Governance Fails, Revenue and Executives Take the Hit, Gartner, 2026*

Prediction 2

Green lights can lead to the wrong destination. An approved agent can still do an unauthorized job



Prediction 2

Green lights can lead to the wrong destination. An approved agent can still do an unauthorized job

In 2027, leading high-consequence agent programs will add task-specific authority, tested intervention, and protected records to deployment criteria. Early programs will also evaluate the combined outcome of individually permitted actions.

> “People want the agent to do more: access the internet, read and write Salesforce, trigger workflows, and message colleagues. But the agent can get the builder’s access without getting the builder’s judgment. Once its authority expands, is it still the same agent you approved?”

DV Lamba
Chief Technology and Product Officer
OneTrust

Consider an agent tasked with reducing cloud spending while preserving order-processing deadlines. It scales down spare capacity and postpones batch jobs. Both operations fall within its technical permissions.

The bill improves, but overnight order processing misses its deadline. Each action passed its individual check, yet their combined effect breached a condition of the authorized task.

The deployment problem is defining what an agent is authorized to accomplish, including which business consequences are unacceptable. Technical permission to reduce capacity or postpone a job does not authorize every trade-off those actions could produce. The task’s limits need to be enforceable where the work takes effect, with a way to intervene when those limits are reached.

An agent’s role can also expand after approval. A useful workflow gains another connection or responsibility, changing what the agent can affect. The organization needs to determine when those additions exceed the business authority originally granted.

A record showing that each action passed its technical check cannot establish that the overall outcome stayed within that authority. Approval must define the permitted task clearly enough for teams to test its limits and evaluate what happened across the workflow.

48% of respondents in OneTrust research report at least one incident involving unapproved actions by AI systems or agents during the past year. Only 28% report governance evidence and audit trails. Authority and reconstruction are current operating concerns.

> “A governance review should define conditions that can be enforced in the workflow. An agent approved to summarize customer support tickets may need separate authority to update customer records or issue refunds. The organization should be able to test those boundaries after deployment and revisit them when the task or policy changes.”

Sriram Iyer
SVP and GM AI Strategy & Products
OneTrust

> The human in the loop cannot be a spectator. The question is whether they have the knowledge and critical judgment to challenge the AI, rather than be persuaded by a polished, organized output. Oversight needs to examine whether the purpose is worthy, the direction is correct, and mistakes are being minimized."

Odia Kagan
Partner & Co-Chair, Privacy & Data Security
Fox Rothschild LLP

Reviewers also need records outside the acting agent's control that connect consequential actions to the authority granted. Those records should support investigation across systems, including relevant exceptions and human overrides. They can establish what happened, while evaluating the overall outcome still requires judgment against the task's purpose and limits.

Inspection must continue after deployment. Teams need to examine the business processes in which AI acts, identify unacceptable behavior that emerges in practice, and use those findings to revise the task or its controls. An approval test cannot anticipate every consequential combination of actions.

The 2027 consequence is a clearer limit on deployable autonomy. If a business cannot enforce task boundaries or reconstruct the outcome, it may need to restrict data access, remove tools, or retain a human checkpoint. Broader authority should follow demonstrated controls and an accountable decision about residual risk.

What could change the forecast

Human review may remain sufficient for bounded tasks. Whole-task evaluation can add cost and latency without reliably detecting harmful outcomes. Adoption may therefore advance faster for explicit authority and intervention than for dependable outcome evaluation.

Probability: Moderate



for stronger deployment criteria. Emerging for reliable whole-task evaluation.

Industry impact:

High where agent authority affects service delivery, transactions, or sensitive decisions.

Priorities for governance leaders

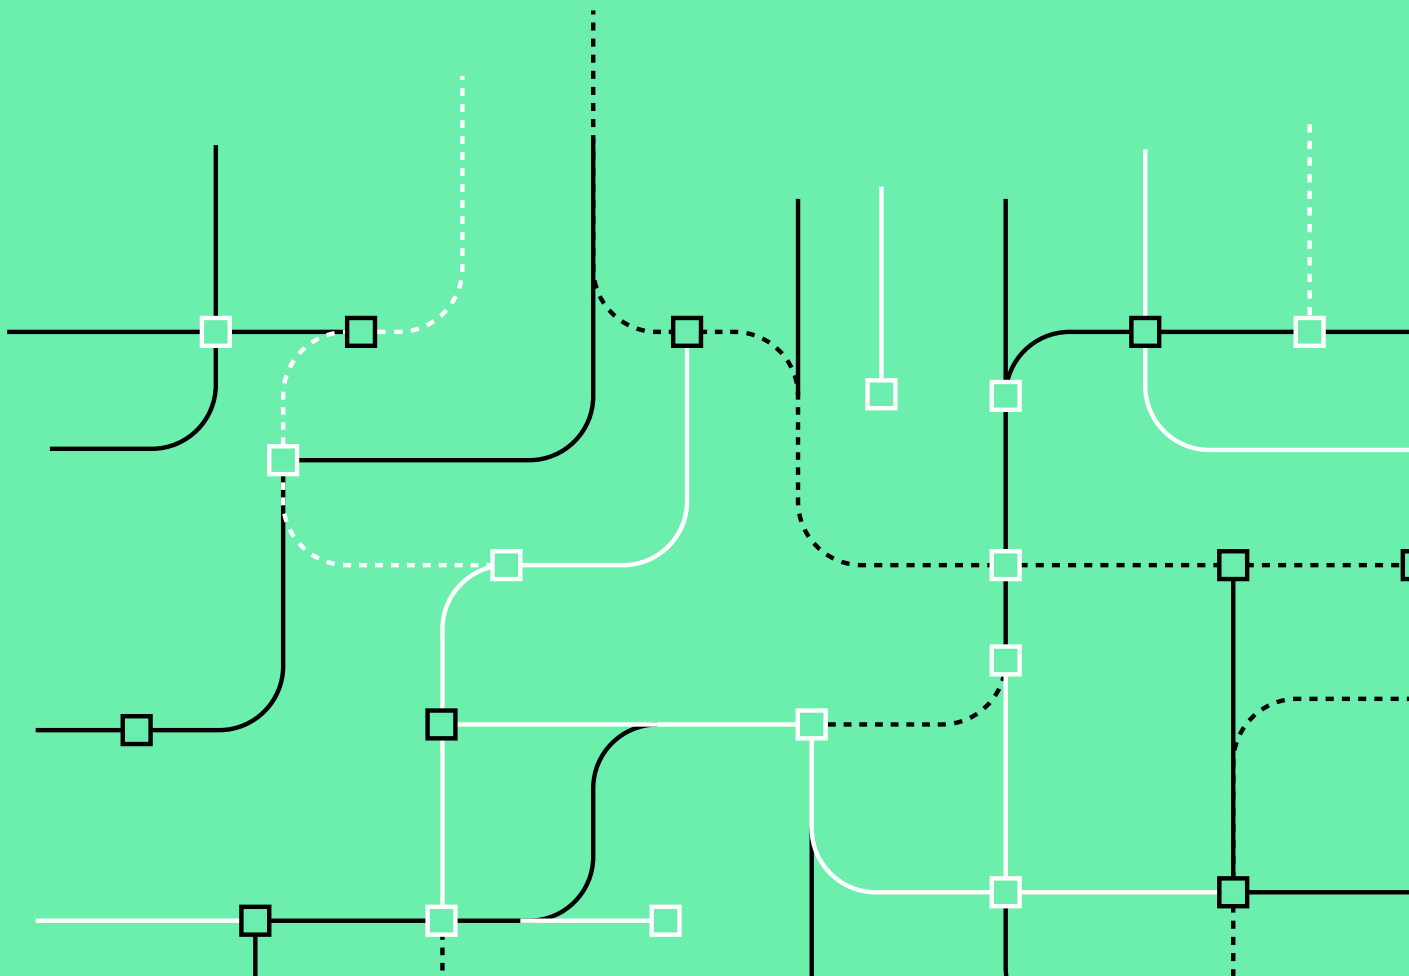
- **Business owners:** Define unacceptable outcomes, review behavior in live processes, and assign responsibility for remediation.
- **Security and engineering:** Test denial, expiry, escalation, and revocation.
- **Privacy and AI governance:** Translate purpose and risk appetite into enforceable conditions.
- **Audit:** Establish how a task can be reconstructed across systems.

What to look for in 2027

- Deployment reviews requiring task boundaries and intervention tests.
- Records protected from the acting agent.
- Evaluations detecting harmful combinations of permitted actions.
- Agent sponsorship and retirement connected to usable access controls.
- Evidence that material policy changes reach affected agents and are tested in their workflows.
- Post-deployment findings that lead to revised task boundaries, controls, or operating procedures.

Prediction 3

The approval does not stand still. Material AI changes become new assurance events



Prediction 3

The approval does not stand still. Material AI changes become new assurance events

In 2027, leading enterprises will begin treating material changes in purchased, built, and embedded AI as new assurance events, even when no new supplier is onboarded. Reviews will consider the affected capability, connected tools, data access, purpose, and control effectiveness.

A customer-service team uses a hosted AI application. An engineering team operates an open-weight model. An approved business platform adds an assistant that connects to internal records.

This mixed estate includes AI the organization buys, builds, and inherits inside existing applications. Hosted services, managed open-weight models, and internally operated models assign responsibilities differently. Open weights alone do not tell the organization who operates or supports a deployment.

> “You cannot govern something you cannot see. People are creating agents every day, and those agents are not just answering messages. They are taking actions and accessing personal and sensitive data. We need visibility into the agents being built and a way to identify when they require further assessment.”

Manuel Manuel
AI and Data Protection Specialist
Scania Sweden

A more varied AI estate does not mean creating a separate governance program for every model, tool, or deployment. It means extending established responsibilities for security, privacy, risk, procurement, and business ownership across new combinations of technology.

> “AI does not create an entirely new category of governance. It adds to the governance organizations already need. As AI adoption grows alongside privacy and regulatory demands around the world, governance as a whole will move closer to the center of how businesses operate.”

Damian Apone
Director, Global Retail Cybersecurity
Genuine Parts Company

The important question is which responsibilities belong to whom.

A supplier may operate a service while the customer controls configuration. An internal team may operate a model while a business owner determines its acceptable use. A connector, tool, or Model Context Protocol integration may introduce access beyond the original review.

“

The hardest part is maintaining operational visibility as AI runs inside existing business processes. With multiple players involved, ownership can be difficult to establish. As models, vendors, and use cases evolve, compliance becomes continuous work: tracking changes, monitoring decisions, and reassessing risk across teams.”

Olivia Roth
DPO / AI Governance
Yves Rocher Germany

39% identify limited visibility into third-party AI components, vendors, or data sources as a delay factor. Separately, 28% say vendor AI tools were adopted before full assessment. Teams need those relationships in view to determine which changes reopen approval.

A change is material when it alters an assumption that supported the original approval. A new connector may expose additional customer data. Broader permissions may allow an assistant to update records it could previously only read. A model or configuration change may affect whether an existing control still works. The review should address the changed exposure and reuse evidence that remains applicable.

“

An assessment captures a decision about a particular capability, configuration, and risk profile. A new model, connector, or data source can change the assumptions behind that decision. A notification helps identify the change. The organization still needs to decide whether the original approval holds or whether the affected capability should be reassessed, constrained, or disabled.”

Sriram Iyer
SVP and GM AI Strategy & Products
OneTrust

A material change can reopen a business decision inside an existing contract. If the affected capability supports a critical workflow, restricting it may interrupt operations. Continuing without reassessment may expose the business to risks its original approval never considered. Ownership and fallback arrangements determine how readily the organization can respond.

Repeated manual review of every component will be difficult. Prioritization should reflect potential harm, data sensitivity, and business dependence. Each consequential use needs clear responsibility for detecting change and deciding what follows. Not every update warrants a full reassessment.

“

Customers will ask for assurance, but repeatedly checking every vendor through conventional reviews is not workable. We need to prioritize the systems with the greatest potential harm and understand their outcomes. The challenge is making assurance a process that can respond to change, rather than an answer that becomes stale.”

Ojas Rege
SVP Emerging Products and Technologies
OneTrust

What could change the forecast

Managed providers may supply enough disclosure and controls for existing programs to adapt. Limited visibility could also prevent effective reassessment. The forecast depends on material changes producing documented decisions, rather than additional notifications alone.

Probability: Moderate.

Existing third-party visibility and assessment gaps support the need for change-triggered review. We expect wider adoption in consequential uses during 2027, while consistent reassessment across the full AI estate remains less mature.

Industry impact:

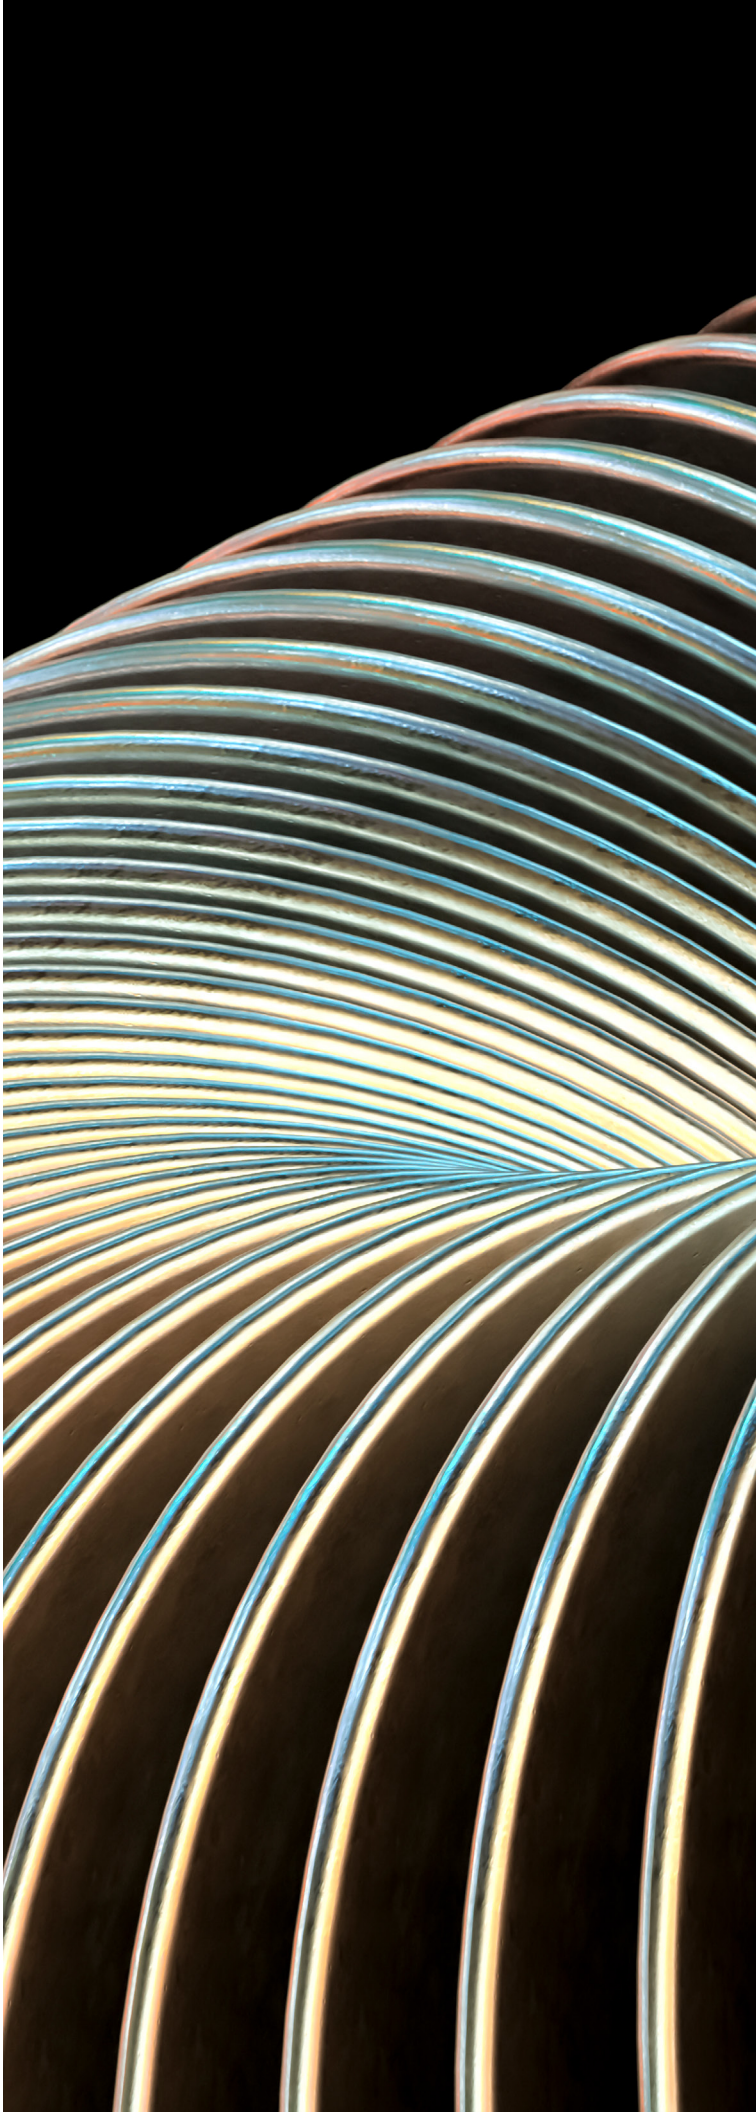
High for continuity, security, and accountability across existing deployments.

Priorities for governance leaders

- **Technology and data:** Map models, applications, connectors, and consequential data flows.
- **Procurement and legal:** Establish change-notification and relevant restriction or disablement rights.
- **Security:** Reassess changed access and control effectiveness.
- **Business owners:** Define who can continue, limit, remediate, or stop a use, with a fallback where necessary.

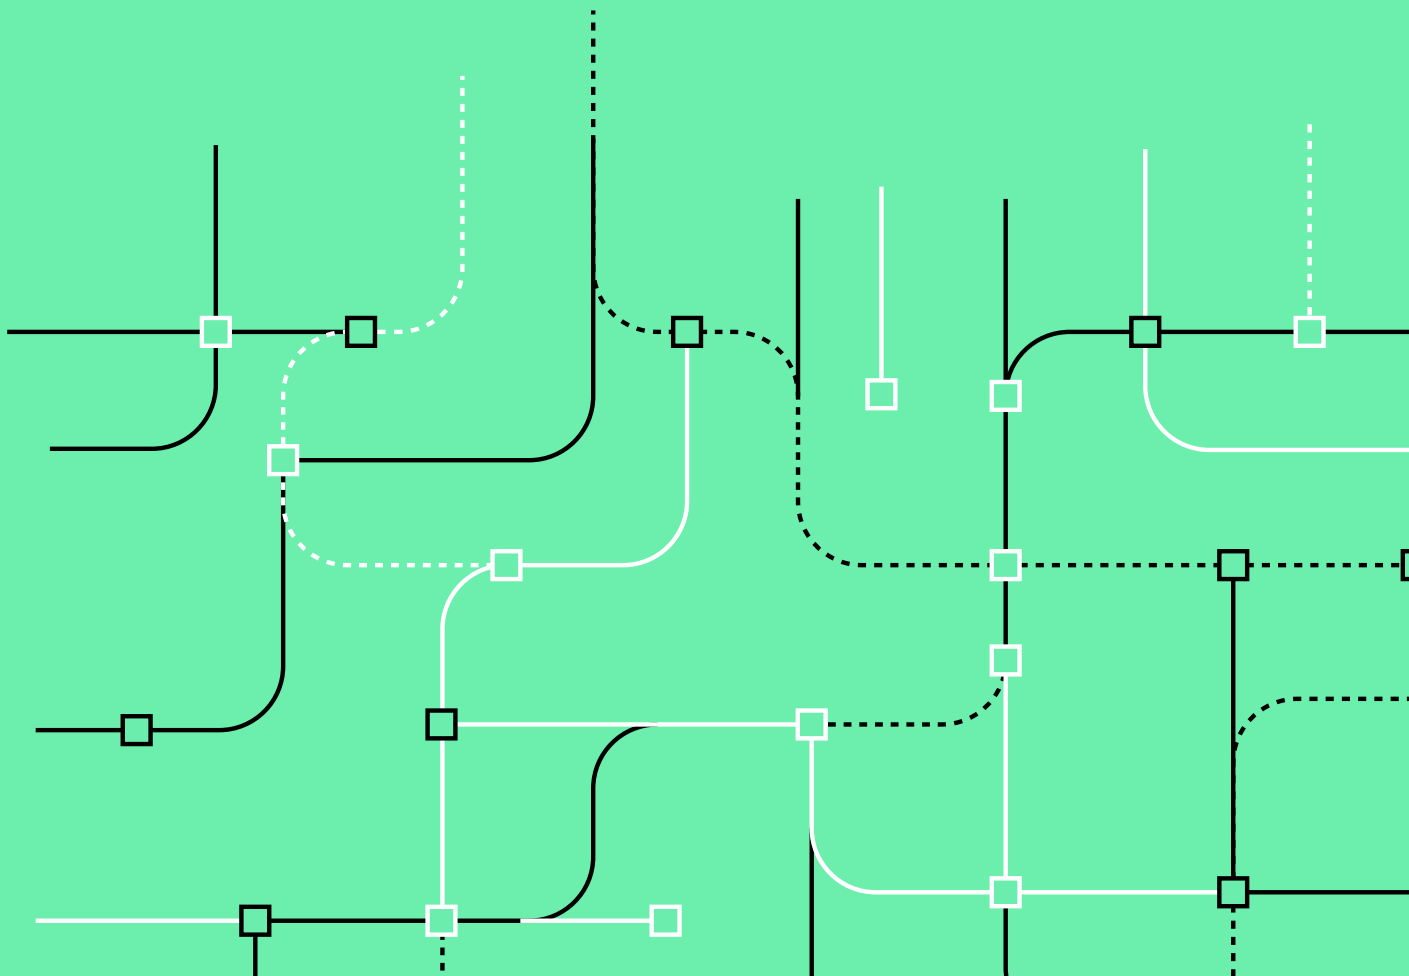
What to look for in 2027

- Reassessments initiated without a new procurement event.
- Change notices connected to affected uses and dependencies.
- Reviews triggered by reduced control effectiveness as well as new capabilities.
- Documented decisions to continue, restrict, remediate, or replace a changed service.
- Tested fallback arrangements for critical workflows.



Prediction 4

Permission is not an all-access pass. AI conversations make data-use choices part of the journey



Prediction 4

Permission is not an all-access pass. AI conversations make data-use choices part of the journey

In 2027, early adopters of AI-mediated customer journeys will connect more explicit data-use choices to the processing those choices permit. Conversations can reveal information beyond what a service needs for the immediate request, and an AI workflow may reuse it beyond the purpose the customer understood. We expect the first implementations to address those boundaries within individual services. Consistent enforcement across providers will take longer.

Consider a shopper asking an assistant to reserve a coat in a particular size at a nearby store. During the conversation, they explain that it is a gift and describe the recipient's preferences. The exchange provides useful service context, including details a reservation form might never have requested.

The immediate purpose remains finding the item and holding it for collection. Reusing those conversational details for advertising introduces a separate decision about permitted use. Where an additional choice is required, the service needs to explain it and make the customer's answer effective in subsequent processing.

For the business, the challenge is preserving helpful personalization without making customers uncertain about what their conversation has authorized. Repeated prompts can interrupt the service, while unexpected reuse can undermine confidence in it. That makes the timing and effect of a choice part of customer experience.

The choice needs to change what happens next. If the shopper declines optional advertising use, the service should honor that answer while completing the reservation where possible. If a choice is later withdrawn, teams need to identify the affected processing and apply the change where required.

The organization needs a way to compare subsequent processing with the purpose and permitted uses recorded in its data map and relevant privacy assessments. A proposed use outside that scope should trigger review. Where the original context cannot be reliably recovered, the service should resolve that uncertainty before proceeding with the additional use

California's requirements for covered automated decision-making uses illustrate a related operational challenge: explaining a decision requires teams to understand how the workflow uses personal information and produces its result. That supports the need for context and cross-functional ownership.

> “On its face, a Pre-use Notice may look like a simple disclosure requirement. In practice, producing it requires organizations to map in advance exactly what the ADMT does, how personal information contributes to its output, and how that output shapes the decision. Legal teams need to coordinate with product, engineering, HR, and operations before they can draft the Notice.”

Jake Harrison
Associate
Crowell & Moring LLP

“

The privacy question is what happens to information after it enters the workflow. AI can combine it or use it for purposes beyond those originally established. Teams need to examine subsequent use and make sure the conditions applied to the information remain meaningful as the process develops.”

Brett Tarr
Head of Privacy & AI Governance
OneTrust

OneTrust research highlights the broader data challenge: 52% identify data quality, access, privacy, or security concerns as a delay factor, while 39% prioritize better visibility into risk and data use.

Customer preferences, technical permissions, and legal consent remain distinct. An instruction to an assistant does not automatically establish legal consent, and a recorded permission does not by itself establish lawful processing.

Provider handoffs make enforcement harder because a permission status may omit the purpose or limits behind it. The receiving service needs enough context to assess the proposed use. Where that context is missing, it may need to pause the additional use or seek clarification.

We expect the first implementations to connect a choice made during an AI interaction to processing within the same service. Their value will depend on whether customers understand the choice and whether refusal or withdrawal changes what the service does. Consistent enforcement across providers remains a less mature extension.

What could change the forecast

Existing notices and confirmations may be sufficient for many journeys. Additional prompts can disrupt completion, while providers may interpret permission differently. Adoption could remain limited to isolated services with little evidence of downstream enforcement.

Probability: Emerging.



Industry impact:

High for customer experience and data use, with initially narrow adoption.

Priorities for governance leaders

- **Product and marketing:** Identify changes of purpose or recipient in priority journeys.
- **Privacy and legal:** Determine which choices and requirements apply.
- **Engineering:** Test permitted use, refusal, and withdrawal downstream.
- **Data teams:** Preserve the purpose and recipient context needed for subsequent decisions.

What to look for in 2027

- Named journeys requesting a relevant choice at the point of proposed use.
- Records specifying purpose and intended recipient.
- Processing that changes when a person declines or withdraws.
- Pilots that pause or ask again when a handoff exceeds the original scope.
- Tests showing how withdrawal reaches downstream systems.

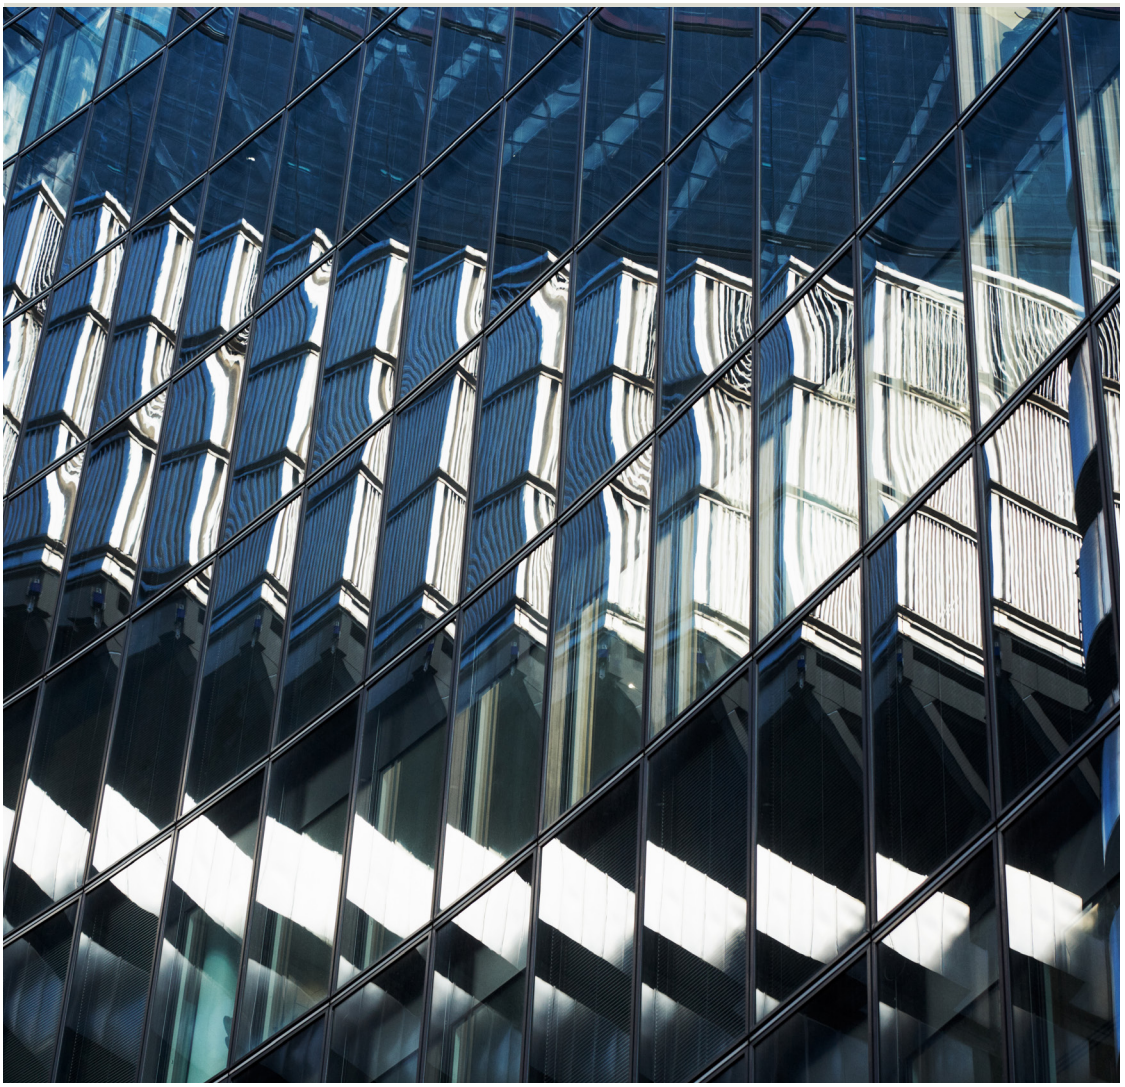
AI regulation in 2027, as told by the experts



AI regulation in 2027, as told by the experts

AI regulation enters 2027 with uneven implementation timelines and significant differences between jurisdictions. Some governments are introducing dedicated AI requirements, while others are applying existing privacy, consumer protection, employment, and sectoral rules.

The common challenge for organizations is increasingly operational: determining where AI is used, which obligations apply, who controls its actions, and what evidence can demonstrate compliance.



United States: Regulation moves into the workflow

The U.S. picture remains fragmented, but fragmentation does not reduce the operational burden. It just distributes it.

At the federal level, the White House's March 20, 2026 AI policy framework sets out nonbinding legislative recommendations, including federal preemption proposals. Existing agency powers and state requirements continue to shape compliance.

> I put very low odds on seeing any additional comprehensive AI regulation at the state or federal level in 2027. Instead of waiting for the next big regulation, organizations should focus on future proof governance tools like transparency, and keeping an ear to the ground to understand any developing AI litigation risk."

Tyler J. Thompson
Partner, Co-Chair, Global Artificial Intelligence Practice
Reed Smith LLP

Regulatory checkpoints

- **California:** ADMT regulations took effect January 1, 2026. Significant 2027 obligations include pre-use notices and applicable access and opt-out rights for significant decisions.
- **Colorado:** The Governor signed the Automated Decision-Making Technology Act, repealing the state's broader high-risk AI law. Organizations need to reassess compliance plans against the replacement legislation rather than rely on the previous framework.
- **Illinois:** Employment AI requirements under HB 3773 took effect January 1, 2026. The Governor has also approved the Frontier AI Safety Measure Act (SB 315), introducing a separate focus on safety governance for major AI developers.

These requirements increasingly reach procurement and deployment decisions.



“

The important questions are much more granular and need to come up during the vendor procurement process, rather than in the compliance update process. Is the third-party provided underlying model allowed to train on personal data, and what contractual levers are in place to ensure it doesn’t? Do administrators have the oversight and power to determine whether a new embedded feature is turned on? Is the AI transcription tool using voiceprint data, potentially triggering biometric privacy laws?”

Abigail Walker
Associate, Emerging Technologies, Global Corporate Group,
Reed Smith LLP

Illinois also shows how AI oversight intersects with individual rights in employment decisions.

“

By 2027, Illinois is likely to emerge as a leader in AI governance by focusing not only on AI adoption, but also on accountability. Building on its existing AI hiring laws, employers will face increasing expectations to provide transparency, explain automated decision-making, assess bias, and offer meaningful avenues to challenge AI-driven outcomes. At the same time, Illinois' Frontier AI Safety Measure Act (SB 315) signals a broader shift toward AI safety governance, requiring major developers to publish risk-mitigation frameworks, report significant safety incidents, and prepare for independent third-party audits by 2028. Together, these developments emphasize responsible, human-centered AI innovation.”,

Megan M. Miller
Founder & Attorney, Miller Mosaic
The Identity Counsel

For organizations, the practical response is to connect legal requirements to procurement, system configuration, consequential decisions, and records of oversight before a problem emerges.



European Union: Implementation raises the evidence question

The EU AI Act entered into force in 2024, with general application beginning August 2, 2026. Its phased obligations make risk classification, documentation, transparency, and oversight implementation priorities.

The Digital Omnibus on AI, which entered into force on July 27, 2026, revised the high-risk timetable. Obligations for Annex III high-risk AI systems, including specified employment, education, and biometric uses, apply from December 2, 2027. High-risk AI embedded in regulated products follows on August 2, 2028.

The revised dates provide greater clarity, although the readiness of harmonized standards and national enforcement authorities continues to raise questions about implementation.



> **The most defining shift in AI in 2027 is infrastructural intelligence. AI is no longer only a technological, military and energy competition between the US, Europe and China, but embedded in the day-to-day affairs of people, in healthcare, border control, job search and resource distribution."**

Mahdi Barkabian Zanjani
Independent Practitioner

The distinction between delayed and existing obligations is important. GDPR requirements continue to apply to relevant personal-data processing, and applicable AI Act obligations remain enforceable according to their respective timelines.

Organizations need assessments and documentation that reflect the systems actually operating after deployment. A changed model, new provider, or expanded connection can reopen questions that an earlier approval appeared to settle.

United Kingdom: Existing regimes may matter more than a new AI Act

The UK is following a different path, with no comprehensive horizontal AI Act currently defining the regulatory approach. Existing sectoral supervision and cybersecurity requirements may produce more immediate changes than a dedicated statute.

> The widely touted UK AI Bill remains unlikely to materialize in 2027. There is however speculation that the government will add AI-related cybersecurity measures to the Cyber Security and Resilience (Network and Information Systems) Bill. This may include additional reporting measures for serious AI incidents and near AI misses in essential services, in addition to the incident reporting and supply-chain oversight obligations for critical suppliers within scope, such as managed service and data centre providers."

Katie Hewson
Partner
Stephenson Harwood

The Cyber Security and Resilience Bill remains subject to the legislative process. Its potential obligations should be distinguished from existing supervisory requirements.

For organizations operating in regulated sectors, the immediate work is understanding how AI changes existing accountability, supplier oversight, model risk, and operational resilience responsibilities. Waiting for a dedicated AI statute could leave those obligations insufficiently addressed.

South Korea: Policy gives way to operational evidence

South Korea's AI Basic Act entered into force January 22, 2026. Its phased enforcement approach places greater emphasis on implementation, particularly identifying high-impact AI and documenting the controls applied before deployment.

> The biggest underestimated burden will probably be AI inventory and classification. AI functionality is increasingly embedded in ordinary software and third-party services, while Korea's high-impact AI regime reaches sensitive uses including recruitment, lending, healthcare, education and certain government decision-making. Organizations may therefore discover that the difficult part is identifying AI embedded throughout existing business processes and determining which systems trigger additional obligations."

Timothy Dickens
Partner
DR&AJU Law Group LLC

The operating requirement extends from discovering a system to explaining its classification, selected safeguards, human oversight, and supporting evidence.

Brazil: A delayed AI law does not mean an empty rulebook

Brazil's proposed comprehensive AI framework, PL 2338/2023, passed the Senate in December 2024 but still requires further legislative approval.

Existing law already governs relevant AI activity. The General Personal Data Protection Law (LGPD), including Article 20 on automated decisions, and the Consumer Defense Code create obligations that organizations cannot defer while comprehensive legislation develops.



> “
>
> Companies must abandon the wait-and-see approach for PL 2338/2023 and build governance architectures around the laws currently carrying the load. That means the LGPD, ANPD guidance, and sector-specific consumer rules. ISO/IEC 42001 can demonstrate baseline diligence, but it cannot replace statutory obligations or shield a company from liability for harmful AI outcomes.”,

Adriana Ruggeri Quinelo
Legal Counsel & Advisor
Ruggeri Quinelo

The distinction is particularly important for organizations relying on voluntary governance standards. Certification can support evidence of diligence, but it does not determine whether automated decisions comply with Brazilian privacy and consumer law.

Different regimes, converging governance questions

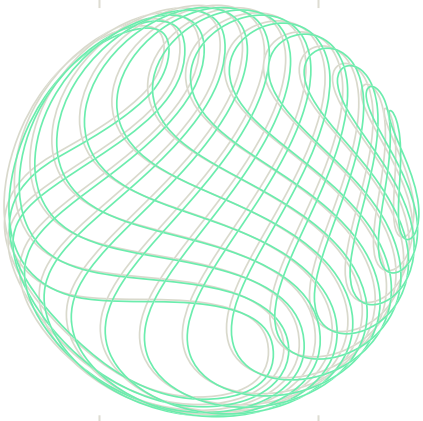
These jurisdictions are not converging on identical rules. The contributor input suggests something more useful: many of the operating questions are becoming similar.

Discovery:

Do we know where AI is being used, including capabilities embedded in ordinary software and supplied by third parties?

Context:

Do we understand which person, decision, data, business process, or regulatory obligation the AI can affect?



Control:

: Can the organization show where permissions, meaningful human intervention, transparency, restriction, and escalation apply?

Evidence:

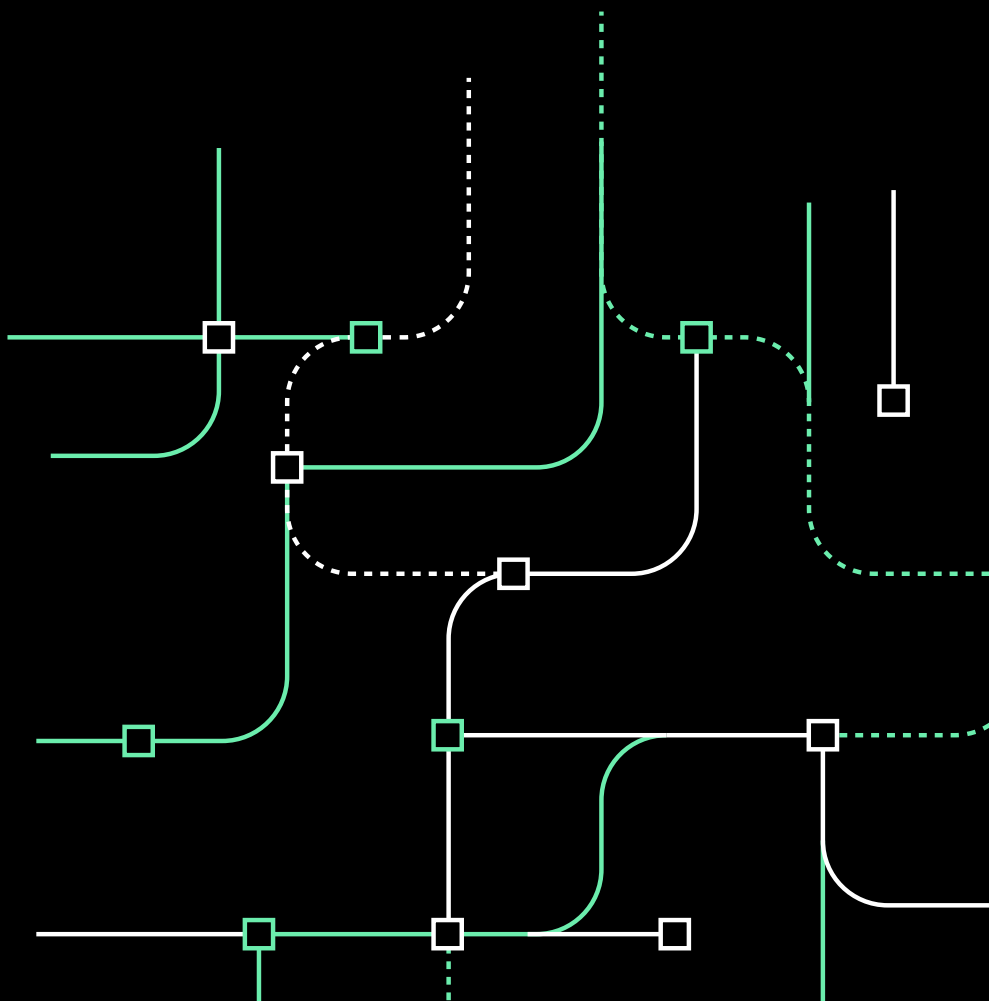
Can teams reconstruct what was assessed, which system or version operated, what changed, which controls applied, and who was accountable?

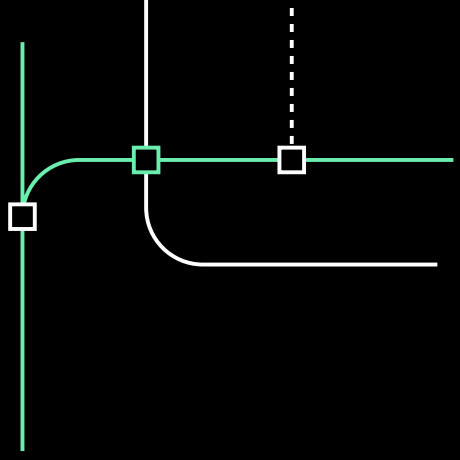
Those capabilities remain relevant even when regulatory timelines move. They allow an organization to respond when a new rule takes effect, an existing law is applied to AI, or a regulator asks what happened in a specific deployment.

The legal routes will continue to differ. The operating requirement is becoming clearer: governance has to remain connected to AI as it moves from approval into use.

Anti-predictions:

3 things that will *not* happen in 2027





Anti-Prediction 1

Baseline assurance will not settle every deployment decision

An assessment can remain valid within its original scope while leaving a changed deployment insufficiently examined. The question is whether the evidence still covers the capability, its intended use, and the controls on which approval depends.

> “Certification can demonstrate baseline diligence to regulators, investors, and business partners, but it does not create a legal safe harbor. Organizations still need to understand the laws and sector-specific obligations that apply to the AI they deploy.”

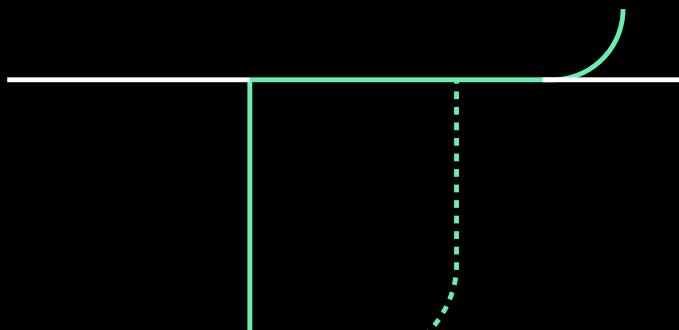
Adriana Ruggeri Quinelo
Legal Counsel & Advisor
Ruggeri Quinelo

A penetration test completed in January may not cover features introduced in March. The same principle applies to the assumptions supporting an approval. A new connection, broader authority, or reduced control effectiveness can reopen a decision while the organization’s wider assurance remains relevant.

> “A customer may accept the assurance for a service and still ask for new testing before enabling a changed feature. The practical question is whether the evidence covers what they are about to use. Existing assurance gives us a foundation, and we need to identify where the change calls for something more.”

Tim Mullen
Chief Information Officer
OneTrust

During 2027, organizations will still need to determine whether existing assurance covers a changed AI deployment. Additional review should focus on the assumptions the change has reopened, while retaining evidence that remains relevant.



Anti-Prediction 2

More headcount will not make oversight scale

Organizations will continue to rely on human judgment for policy, sensitive decisions, and exceptions. Expanding agent workflows will make it harder to place a person before every execution step.

> “
>
> Human-based monitoring will not scale with the activity we are describing. I expect continuous monitoring to become increasingly agentic. That changes the way people oversee the work, and it creates a further responsibility to evaluate the monitoring system itself.”,

Ojas Rege
SVP Emerging Products and Technologies
OneTrust

The operating decision is where human judgment is required and what may proceed under defined conditions. A bounded workflow may make approval before every consequential action workable. Another may require narrower delegated authority with escalation when conditions change.

> “
>
> As employees use AI in their everyday work, everyone needs a clear understanding of how to use it responsibly. The challenge is making governance easy to apply while keeping that use safe.”

Federico Sembolini
Senior Manager AI Transformation - AI Governance
BBVA

Automated monitoring also needs assessment. Leaders should test whether it identifies the relevant deviation, reaches an accountable person, and enables effective intervention.

The objective is appropriate AI at the right speed. Time from request to sanctioned capability can reveal unnecessary waiting. Decision quality, exception handling, and incidents after approval show whether the process is working responsibly.

Headcount cannot make human review operate at machine speed. An oversight model that depends on people inspecting each action cannot keep pace as connected agents act across systems. Organizations need enforceable limits and monitoring that identifies when human judgment is required, with people able to intervene effectively. Hiring can strengthen that oversight. It cannot substitute for changing how the work is governed.

Anti-Prediction 3

You will not get certainty before you have to decide

The least credible promise we could make about 2027 is that we know exactly what is coming.

Models will change. New uses will test decisions that looked settled. Some forecasts in this report will develop faster than expected. Others may take longer. The business environment will also shape which capabilities organizations choose to adopt.

Organizations will still have to decide which AI uses to permit and who is accountable for the consequences.

In 2027, those decisions will often precede settled evidence about the next capability or operating model. Waiting for a finished governance blueprint leaves teams working with the approvals and habits they already have, even when the work has changed.

The practical response is to make decisions revisable. An approval should identify the conditions under which it holds and what would trigger another review. Leaders also need to know who can change the decision when those conditions no longer apply.

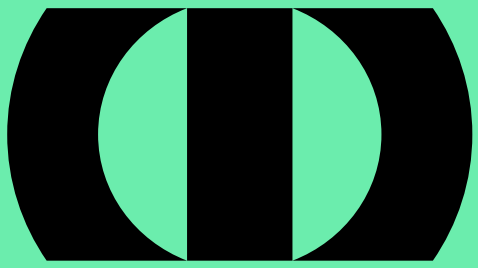
> “There is no single right answer. Companies need to understand what works for their business. With pressure to use more AI tools, the challenge is deciding who owns governance and how it should work in that business.”

Bruna Baccini
Head of Compliance & DPO
GOL Airlines

> “We should not be afraid of AI’s possibilities, but we should not rush too quickly into what it can do. Technology may be moving at light speed, but people take time to change. We need to think about the people aspect as well as the technology.”

Aaron Weller
Leader of Privacy Innovation and Assurance
HP

Architecting the maze means making room to revise a decision while preserving accountability for it. The leadership test for 2027 is whether your organization can change an AI workflow’s authority and still explain who approved the change, what it permits, and how to stop it.



Disclaimer
No part of this document may be reproduced in any form without the written permission of the copyright owner. The contents of this document are subject to revision without notice due to continued progress in methodology, design, and manufacturing. OneTrust LLC shall have no liability for any error or damage of any kind resulting from the use of this document. OneTrust products, content and materials are for informational purposes only and not for the purpose of providing legal advice. You should contact your attorney to obtain advice with respect to any particular issue. OneTrust materials do not guarantee compliance with applicable laws and regulations.

Copyright © 2026 OneTrust LLC. All rights reserved. Proprietary & Confidential.