



State of AI Governance

2026

The Policy-to-Proof Gap in
AI-Assisted Communications

Research conducted by



About This Benchmark

Regulated organizations have seen this pattern before. A new tool enters daily work before rules are fully set. Business teams move quickly, and Compliance is expected to answer when questions arise.


AI follows that pattern, but it creates a harder evidence problem. AI-assisted work may involve prompts, generated content, human edits, reviews, approvals, and final communications. If those pieces are not connected, it becomes harder to show what happened.

Compliance is everyone's responsibility, and it requires clear accountability. Whether the issue is money laundering, insider trading, recordkeeping, or AI, responsibility is shared across Compliance, Legal, IT, and business teams. Regulators expect senior leaders to oversee these programs, fund them appropriately, and demonstrate that controls work.

To measure readiness, Arctera commissioned an independent study of 500 Compliance, Legal, Risk, Regulatory, and IT decision makers in Finance, Healthcare, and Energy. Respondents represented organizations with at least 500 employees and \$50 million in annual revenue.

The study found that 85% of regulated organizations already use AI, creating an urgent need for evidence readiness. Although 80% have policies, far fewer can defend their AI-assisted actions. Teams now need to show how AI was used, what it helped create, and whether that work can be reconstructed if questioned.

This benchmark centers on three readiness questions: where AI enters regulated work, where the evidence trail can break, and whether the organization can reconstruct what happened.



AI readiness requires more than policy coverage. Organizations need to track where AI is used, maintain a connected evidence trail, and be prepared to defend their actions to auditors and regulators.

That story reveals itself in six signals. They show where AI governance is advancing, where evidence readiness is falling behind, and what leading organizations are doing to close the policy-to-proof gap.

Research conducted by



The Six Signals of the Policy-to-Proof Gap

The benchmark points to a clear policy-to-proof gap. Regulated organizations are accelerating AI adoption and formalizing governance policies. These advances do not, on their own, create a defensible record. The real test comes when organizations must prove the record: where AI was used, what it produced, who reviewed or changed it, and whether the resulting record can be reconstructed.

As AI-assisted work moves through regulated workflows, that record becomes harder to preserve. Prompts, edits, approvals, and final records each become part of the evidence trail. For governance to be defensible, that evidence has to remain connected.

Those findings unfold across six signals, moving from AI adoption and accountability to evidence capture, reconstruction, channel coverage, and archive readiness.



Together, the signals show that AI governance readiness depends on connected evidence, not policy coverage alone.

From Policy Confidence to Evidence Readiness

Policies matter. They define expected use, review points, and guardrails. They do not show what happened inside a workflow.

In regulated work, organizations need the evidence behind the work itself: source context, AI assistance, human review, approvals, and the final communication. When those pieces are disconnected, teams may know AI was used but struggle to explain how it influenced a communication or decision.

As AI-assisted work enters the corporate record, evidence readiness depends on capturing the full lifecycle: what AI produced, who reviewed it, what changed, and how the record was preserved.

Where the Record Starts to Break Down

This derived view from the benchmark captures the central readiness problem in one line. Many organizations can capture AI activity at the start. Readiness drops as the work moves into retention, review, and investigation. By the final stage, only 18% begin with most of the necessary data ready.



Research conducted by



At the point of scrutiny, readiness is weakest. Organizations may have pieces of the AI record, but few can connect the full path from AI output to the final record.

The six governance signals that follow examine where that path breaks, beginning with the first and most immediate question: where is AI already entering regulated work?

AI Has Entered the Regulated Record

AI-assisted work is already present across communications, summaries, recommendations, records, and decision-support workflows. A financial advisor may use it to draft a client follow-up. A healthcare team may summarize a case history before review. An energy or utilities team may prepare an operational update after an incident. A compliance analyst may summarize communications before an investigation.

These use cases can improve speed. They also create evidence that must stay connected from source input to final record.

Visibility is the first governance test. Teams need to see where AI is used, what it produces, and how that work becomes part of the regulated record. When that view is incomplete, organizations may know AI was involved but struggle to explain how it shaped a communication, recommendation, or decision.

Regulators are already applying existing compliance expectations to AI. In financial services, FINRA has said securities rules apply to GenAI and similar technologies, including embedded AI features in third-party products. Recent SEC “AI-washing” actions also reinforce the need for accurate, supportable statements about AI use.


What the Data Shows

The adoption finding itself is not the surprise. Regulated organizations were always going to use AI where it could speed drafting, summarization, review, and decision support. The more important finding is how quickly that routine use becomes a recordkeeping and evidence problem.

Compliance teams already know this pattern from communications supervision. Risk often becomes visible when business activity moves faster than governed channels can support. AI extends that challenge because the risk may begin before the final communication is sent, in the work that shaped it.

In the benchmark, 78% of organizations report widespread AI adoption in regulated workflows, and nearly as many expect AI-related communications risk to rise over the next two years.

Adoption Is Already Inside Regulated Workflows



AI Communications Risk Is Expected to Rise

78%
expect risk to increase

Increase expected

Research conducted by




The chart makes the baseline clear: AI is already inside the work Compliance is expected to govern. The sharper question is where that risk shows up first.

Where AI Risk Shows Up

Respondents pointed to two practical risks concerns: content that can spread before review, including inaccurate, misleading, or non-compliant content produced at scale (24%), and sensitive information that can move through prompts, outputs, and downstream records (22%).

Once AI-assisted work enters the regulated record, the next question is accountability: who is expected to answer, and do they have access to the evidence needed to defend the record?



These concerns affect the integrity and defensibility of the record. Content integrity is the first concern because AI can help flawed, unsupported, or non-compliant language move faster.

Data protection follows the same path. Prompts and outputs can carry sensitive context into AI tools, summaries, collaboration spaces, and final communications.

The findings confirm that concerns about AI risk are well-founded. AI is already more pervasive in regulated work than many organizations may be prepared to govern. Organizations need visibility into where AI is used, what it creates, and what sensitive context may move through it.

The next question is accountability: who is expected to respond, and do they have the evidence to explain what happened?

Research conducted by

Compliance Carries the Response Burden


AI risk naturally pulls Compliance into the response. Compliance is expected to understand exposure, oversee controls, and support the organization when questions arise. AI touches many of the areas Compliance already watches: communications, records, supervision, investigations, privacy, and regulatory response.

The pressure point is evidence. The records needed to explain AI-assisted work rarely sit in one place. They may live in business systems, collaboration tools, AI tools, archives, records systems, and third-party providers.

That creates the ownership gap. Compliance may be expected to explain what happened, while the evidence often sits with the teams that manage the systems where AI-assisted work happens.

The AI Ownership Gap

In the benchmark, 60% of organizations assign primary accountability for AI-governed communications to Compliance.




That finding confirms Compliance's central role. It also shows where pressure can build. When questions arise, Compliance may need to locate the supporting evidence path, from AI activity and source material to review decisions, retention actions, archive records, and downstream communications.

Dedicated or shared governance makes that response stronger. Organizations with dedicated AI governance report higher levels of monitoring and logging (87%), automated archiving (67%), and audit-trail readiness (90%). Shared governance also performs above Compliance-only ownership across the same controls.

Research conducted by

Clear Accountability Strengthens Technical Controls

UP TO **3x** Higher Readiness



The control gap is significant. Where Compliance is the sole named owner, readiness is lower across monitoring and logging (62%), automated archiving (40%), and audit-trail readiness (69%).

The implication is practical: Compliance needs clearer access to evidence, defined response authority, and shared ownership with the teams that manage the systems where AI-assisted work happens.

The next signal turns to policy and evidence: written rules can clarify expected use, but proof still depends on the technical evidence behind the work.

Research conducted by



Policy Is Not Proof


Regulated organizations have made real progress on AI governance. Many have formal policies, training, human review expectations, and named owners. Those controls are important because they define how AI should be used.

The harder question comes later: what actually happened?

A policy can say which tools are approved, what data should not be entered, and when review is required. But once AI-assisted work moves through daily workflows, proof has to come from the systems where the work happened.

Technical Capability is the Evidence Layer

More than half of organizations report both policy and technical capability for AI communications (54%). The rest are split across less defensible states: policy without technical capability (14%), technical capability without formal policy (14%), or neither layer (18%).




The challenge is connecting policy to evidence. Approved AI use cases need to align with the systems where work actually happens, so teams can capture AI activity, retain generated content, document review steps, and make final records searchable when questions arise.

Research conducted by

Policy Maturity is Not Proof of Readiness

Formal AI policies are common, but the supporting controls narrow quickly. Four in five organizations have a formal AI policy, but fewer than three in ten have the core evidence controls needed to prove AI activity.

The decline appears across the controls that make proof possible: 55% log AI prompts and outputs, 29% automatically retain or archive AI-generated content, and 29% have the full set of core evidence controls in place.




Defensibility is the test. Policies define expected use, but evidence controls show whether those expectations were followed. Without prompt and output logging, automated retention, and connected records, organizations may be able to point to a policy while still struggling to prove the path of AI-assisted work.

Research conducted by

Mature AI Programs Still Have Evidence Control Gaps

Even mature AI programs can have evidence gaps. A program may include documented procedures, review steps, and risk oversight while still lacking the controls needed to produce a complete record.



The remaining gaps are operational. Mature programs still report missing keyword or policy detection (59%), automated retention or archiving gaps (52%), risk scoring or classification gaps (46%), prompt and output logging gaps (25%), and reconstruction readiness gaps (15%).

The takeaway is straightforward: policy maturity does not equal evidence readiness. To defend AI-assisted communications and decisions, organizations need the record of use, review, retention, and reconstruction to stay connected.

The next signal looks at what happens when teams detect AI-related risk but still have to reconstruct the evidence behind it.

Research conducted by

Detection Confidence Is Not Evidence Readiness

The first hour of an inquiry is where AI governance becomes operational.

A monitoring tool may flag a potential issue. A manager may escalate a questionable AI-assisted communication. A regulator may ask for supporting records. Legal may need to understand how a decision was reached. In each case, the signal is only the starting point. The team still has to produce the evidence behind it.

Detection identifies the issue. Reconstruction explains it. When evidence is disconnected, teams spend critical time finding records before they can assess what happened.

The First Hour Disconnect

Many organizations are confident they can spot AI-related risk. The breakdown starts when they have to assemble the evidence behind it. Seventy-two percent are confident in detecting AI-related risk, while only 18% start investigations with most of the necessary data ready.



That gap exposes the difference between knowing where to look and having the record ready. Teams may still need to collect the message, find the related chat, request the AI log, identify the source document, confirm the review decision, check the archive, reconcile timestamps, and determine which version became final.

Without a connected evidence path, the first hour is spent chasing the record instead of understanding the risk.

Research conducted by


The Reconciliation Black Hole

Manual data collection and reconciliation consume critical investigation time.

57%

spend at least a quarter of investigation time gathering or reconciling data before analysis can begin.

Share of Total Investigation Time



Fifty-seven percent of organizations spend at least a quarter of investigation time gathering and reconciling data across systems. That time is more than administrative overhead. Each manual step can introduce missed records, inconsistent exports, incomplete context, unclear ownership, or version confusion.

The more time a team spends rebuilding the record, the less time they have to assess the issue, document the response, and defend the decision.

Investigation readiness depends on having connected evidence before questions arise. The goal is to prevent each inquiry from starting as a data search.

The next signal follows the evidence path into daily workflows, where AI-assisted work moves across tools, channels, and records.

Research conducted by



Evidence Breaks Across channels

AI-assisted work rarely stays in one place. A prompt may start in an approved AI tool. The output may move into a document. A revised version may be pasted into an email. A summary may be discussed in chat. A decision may be confirmed in a meeting. A follow-up may shift to mobile messaging or a third-party channel.


Each handoff may be routine. Each one can also separate the final record from the context that shaped it.

Retaining the final communication is not enough. Teams need a connected path back to the source context, AI assistance, human review, and systems where the work took place.

The Biggest Constraints Are Connected-Evidence Issues

The leading constraints are about connection. Teams do not just need to find the record. They need to connect communications, logs, review decisions, retention actions, exports, and custody into a defensible path.

The Biggest Constraints Are Evidence Continuity Issues



The most common barrier is correlating communications across modalities (37%). Respondents also cite difficulty identifying relevant data across channels and systems (27%). These are evidence problems that show up when AI-assisted work spans multiple tools, channels, or records.

When those elements are disconnected, the organization has to rebuild the path manually.

Research conducted by




Blind Spots Begin Outside Governed Channels

This is where the AI governance problem starts to look familiar. Compliance teams already know what happens when business activity moves faster than governed channels can support. AI raises the stakes because the concern is not limited to the final message. It also includes the work behind it.

33%

cite consumer messaging apps on personal devices as a major blind spot

Where Blind Spots Appear




One-third of organizations flag employee use of AI outside approved platforms as a major blind spot (33%). Often, the issue starts with ordinary business activity: a summary is forwarded, a draft is copied, a decision is discussed in chat, a client question moves to mobile, or a vendor adds context.

When governance does not follow the work, evidence can drift outside the record.

Research conducted by

Channel Coverage Still Lags AI Movement

Coverage gaps make the problem harder because AI-assisted communications can appear in channels where monitoring is partial, inconsistent, or missing.



Nearly half of organizations report partial or no coverage for consumer messaging apps on personal devices (46%), and one in three report partial or no coverage for personal email (34%).

This is the off-channel communications problem in a new form. AI-assisted work can move through the same weak spots compliance teams have spent years trying to control: personal messaging, unofficial channels, third-party tools, and conversations that sit outside approved systems.

The difference is that AI adds another layer of evidence. A final communication may be archived while the prompt, source context, AI-generated draft, review path, or channel handoff that shaped it remains outside the governed record.

As AI-assisted work crosses tools and channels, capture, review, retention, and search need to follow the work. Otherwise, teams may preserve the final message without the evidence needed to explain how it was created, reviewed, and approved.

The next signal looks at whether the archive can become an active evidence foundation.

Research conducted by

The Archive Must Become the Evidence Foundation

Organizations already recognize that archived communications have strategic value. The next opportunity is operational: turning stored records into usable evidence when AI-assisted work is questioned.

AI-assisted work can move across channels, systems, and records. That changes what organizations need from the archive. It can no longer function only as a place to store communications or meet retention obligations. It has to help teams locate records, preserve context, apply holds, support review, and produce evidence when questions arise.

A retained communication is not automatically defensible. Teams still need the context around it: how it was created, whether AI influenced it, who reviewed it, which retention rule applied, whether it was placed on hold, and how it can be produced.

The archive supports AI governance when it enables reconstruction, not just retention.

Archives Are Valued, But Not Yet Operationally Ready

Most organizations recognize the value of historical communications data—70% consider their archive a strategic asset. Yet, only 18% begin investigations with most of the necessary data ready.

Two Benchmark Signals Show the Gap Between Archive Ambition and Investigation Reality


70%

view archived communications data as a major or critical strategic asset

Strategic archive value



Investigation readiness



52-point gap
value is ahead of operational readiness

That disconnect defines the opportunity. The archive may contain the record, but teams still need the surrounding context: metadata, review activity, retention decisions, legal holds, case context, and production workflows. Without those connections, records may be preserved but still lack the context needed for defensible use.


For AI-assisted work, that context is critical. The final communication may be available, while the prompt, output, edit history, approval decision, channel movement, or retention path still has to be assembled from somewhere else.

Research conducted by



What Slows the Move to Connected Evidence?

Archive modernization is slowed more by operational complexity than budget. The leading barriers are security reviews and approvals (36%) and data residency or privacy constraints (35%). Budget constraints rank lower (18%)



Those barriers help explain why the disconnect persists. Organizations are trying to turn archives into usable evidence foundations while navigating security requirements, privacy rules, data residency, legacy architecture, integration complexity, internal IT capacity, and change management.


The goal is faster access to the right evidence in the right context. When an inquiry begins, teams need the complete chain behind a specific AI-assisted record without rebuilding it manually from disconnected systems.

Research conducted by

Technology Investment Priorities

Compliance teams are investing in capabilities that shorten the path from AI activity to defensible response. Integrated AI oversight logging and cross-channel search and reconstruction are each cited by 31% of organizations, followed by real-time monitoring and alerting at 28%.

Top Capability Priorities for the Next 2-3 Years



These priorities point to the archive's next role: connecting AI activity to accountability, records to context, and retention to reconstruction.

That is where the archive begins to support risk reduction, investigation speed, and defensibility. It helps organizations preserve the record, understand how it was created, and produce the evidence needed to explain what happened.

The next section moves from the archive to decision evidence: human review is only defensible when the evidence behind the review can be shown.

Research conducted by

Human Review Needs Evidence

Human review is central to AI governance. It brings judgment, context, escalation, and accountability into AI-assisted workflows.

Review only helps when the organization can show what the reviewer saw and how the decision was made. If an AI-assisted communication is challenged, teams have to show what AI generated, what source content informed it, what changed before approval, and what supported the final decision.

A reviewer may have participated, but the organization still needs the record behind the review: the prompt, output, source material, edit history, risk signal, approval activity, and final communication. If those elements are disconnected, teams may be able to show that review occurred while still struggling to show whether the review was complete, reasonable, and based on the right information.

For regulated organizations, the question is whether they can prove what the human reviewed, what changed, and why the final decision was reasonable.

Human Review Needs an Evidence Trail

Human-in-the-loop oversight depends on the technical audit trail captured behind it.

Human Review Is Stronger When the Evidence Trail is Captured Behind the Review



Only one in three organizations can produce a complete evidence trail for human review or approval. That leaves 66% without the full evidence trail behind the review process.


A final approval without context may not show whether the output was accurate, whether sensitive information was exposed, whether the reviewer changed the recommendation, or whether the final record reflects the approved version.

Research conducted by

Audit Readiness Rises with Connected Controls

Audit readiness improves when review is backed by connected controls: prompt/output logging, review and approval history, and automated retention or archiving.

Defensibility Rises When Review Is Backed by a Connected Technical Evidence Chain



Audit readiness increases as connected controls are added. Among organizations with none of the connected controls, 30% say they can quickly produce a defensible audit trail. Among organizations with the full control set, that rises to 94%.

These controls make the review process easier to defend because they show what the AI system produced, what source content informed it, who reviewed the output, what changed, how it was approved, and how the final communication was retained.

This is where review and evidence come together. Human oversight gives the process judgment. Defensibility depends on the evidence captured inside the workflow: logging, review history, retention, archiving, and reconstruction.

The next section looks at organizations that have begun building this operating model.

Research conducted by

What Proof-Ready Organizations Do Differently


Proof-ready organizations treat AI governance as an operating model. They prepare for the moment when AI-assisted work has to be explained.

They identify where AI enters regulated workflows. They define who owns review, escalation, retention, and response. They capture evidence in governed systems. They reduce manual reconstruction by keeping the evidence path connected before an investigation, audit, or regulatory inquiry begins.

In this analysis, proof-ready organizations are defined by a combined readiness profile: cross-channel reconstruction confidence, defensible audit-trail readiness, faster evidence gathering, and advanced compliance automation. This makes the cohort more than a confidence measure. It reflects whether organizations can see AI-assisted activity, preserve the record, explain decisions, and reconstruct what happened.

The Proof-Ready Cohort vs. Everyone Else

Proof-ready organizations show stronger cross-channel reconstruction, governance alignment, audit-trail readiness, and investigation efficiency.



The cohort comparison shows what readiness looks like operationally. Nearly 9 out of 10 proof-ready organization are confident they can reconstruct activity across channels, compared to just 31% of their peers. An even greater gap appears in audit-trail readiness: 94% of proof-ready organizations can quickly produce a defensible audit trail, while only 35% of others.

Research conducted by

These results show why evidence readiness becomes an operating advantage. Proof-ready organizations are better positioned to respond quickly, reduce manual investigation work, explain decisions, and defend the record.

Reported AI Content Concern Declines with Audit-Trail Readiness

Audit-trail preparedness adds another view of exposure. Among organizations only slightly prepared to produce a defensible audit trail, 54% cite inaccurate, misleading, or non-compliant AI content as a major concern. That concern declines at each higher level of preparedness, reaching 14% among organizations that are extremely prepared.



The finding should be read as a relationship, not a causal claim. The data shows lower reported concern alongside higher audit-trail preparedness. Stronger readiness gives teams a better way to supervise, investigate, and defend AI-assisted content when questions arise.

The most capable programs share a common pattern: visibility into AI activity, clear ownership across teams, controls that preserve evidence, and systems that keep the record connected.

When put to practice, these capabilities create the record the organization has to produce: defensible, connected, and ready for review.

The next section turns that model into a practical readiness framework: what evidence teams need to capture from source to production.

Research conducted by

The Defensible AI Record

Proof-ready organizations are defined by their ability to explain what happened when AI-assisted work is questioned. For organizations still building toward that model, the starting point is the record itself.


Readiness begins by defining what must be available for the AI-assisted workflows most likely to be reviewed, investigated, audited, or produced. These may include customer communications, decision inputs, case notes, employee summaries, incident updates, or risk recommendations.

A defensible AI record has three layers: the inputs that shaped the work, the review actions that governed it, and the lifecycle controls that preserved it.

At the input layer, teams need the source content, user prompt, and generated output. At the review layer, they need human edits, review activity, approval decisions, and the final communication or recommendation. Across the lifecycle, they need retention labels, archive records, legal holds, and production history.

From Source to Production

A connected evidence path helps explain, reconstruct, and defend AI-assisted work across inputs, oversight, and lifecycle management.




Together, these elements show how AI-assisted work moved from source material to final record. They help establish what was used, what changed, who approved it, where it was retained, and how it can be produced. The proof chain has three phases.

Research conducted by

Practical Entry Points for Readiness

Organizations do not need to solve every workflow at once. Readiness can begin where AI assistance is visible, material, and likely to create questions.



For each entry point, teams can define the required record, assign ownership, confirm where the data lives, and test whether the record can be assembled without ad hoc collection.

That work creates a practical bridge from policy to proof. It helps organizations prioritize the records that carry the most regulatory, legal, or operational exposure and begin building the controls needed to defend them.

Research conducted by



CONCLUSION

Evidence Defines the Next Phase

AI governance is moving from policy formation to proof.

Policies, approved tools, and acceptable-use guidance remain important. They set expectations for how AI should be used. The benchmark shows the next test is whether organizations can show what happened when AI-assisted work enters communications, decisions, and regulated records.

That requires more than a retained final message. Organizations need the connected record behind it: source context, prompts, outputs, edits, review decisions, channel movement, retention, archive records, legal holds, and production paths.

The organizations best prepared for scrutiny will build that record into daily workflows before questions arise. They will know where AI entered the work, how it shaped the outcome, who reviewed it, and how the final record can be reconstructed.

For regulated teams, the opportunity is clear: move from policy confidence to evidence readiness. That is how organizations gain clarity into AI use, confidence in response, and control over the governed record.

Set the policy. Capture the evidence. Govern the record.

Research conducted by



Technology and research for regulated organizations managing communications, compliance, and digital risk.

Arctera provides data compliance and governance technology for regulated organizations across finance, healthcare, energy, and utilities. Its solutions support communications governance, retention and archiving, surveillance and review, investigation readiness, legal hold, eDiscovery, and connected evidence.

Arctera helps organizations strengthen the foundations for defensible communications governance by connecting communications, content, context, and evidence across systems and workflows. This includes the visibility and technical infrastructure needed to track AI-assisted content through its lifecycle, identify where it enters the business record, monitor movement across channels, retain critical evidence, and reconstruct records for regulatory or investigative response.

Arctera sponsors independent research to benchmark emerging practices and provide data-driven perspectives for compliance, legal, risk, IT, and executive leadership. This report, The Arctera 2026 AI Communications Governance Benchmark, was conducted in partnership with Hanover Research.

To learn more, visit arctera.com.

About Hanover Research

Founded in 2003, Hanover Research is a global research and analytics firm that delivers market intelligence through a unique, fixed-fee model to more than 1,000 clients. Headquartered in Arlington, Virginia, Hanover employs high-caliber market researchers, analysts, and account executives to provide a service that is revolutionary in its combination of flexibility and affordability. Hanover has been named a Top 50 Market Research Firm by the American Marketing Association every year since 2015. To learn more about Hanover Research, visit hanoverresearch.com.

Research conducted by



An independent benchmark of compliance, risk, legal, regulatory, and IT decision-makers across regulated industries.

Arctera commissioned Hanover Research, an independent market research firm, to conduct the quantitative research for this report.

About the Benchmark: This independent quantitative study was conducted in May 2026 by Hanover Research on behalf of Arctera. Surveying 500 qualified compliance, risk, legal, and IT decision-makers across healthcare, finance, and energy/utilities in the Americas, EMEA, and Australia, the report establishes an empirical benchmark for evidence readiness. All AI-specific statistics are based on the core cohort of n=425 organizations actively deploying AI in regulated workflows.

Respondent Criteria

- Full-time employees aged 22 or older.
- Professionals working in compliance, legal, risk, regulatory, or IT functions across Finance, Healthcare, and Energy/Utilities.
- Managers, directors, VPs, C-suite executives, auditors, or specialists verified as active compliance decision-makers or key program influencers.
- Organizations with at least 500 employees and annual revenue of at least \$50 million.

Data Interpretation

- No statistical weighting was applied to the raw dataset.
- Because respondents were recruited through a qualified, pre-vetted industry panel rather than a probability sample of all global regulated organizations, a traditional margin of sampling error is not reported.
- Findings should be interpreted as directional benchmarks for the qualified respondent population rather than precise estimates of the entire global market.
- Percentages and chart values may be rounded for display clarity.
- The research focuses on digital communications archiving, connected evidence stacks, compliance workflows, and inquiry reconstruction. Broader AI monetization, model training methodologies, and general AI business strategy are outside the scope.

Research conducted by



Key Definitions

AI-Assisted Communications: Any electronic communication or business record created, edited, or summarized—fully or in part—by an artificial intelligence (AI) tool or platform.

Evidence Path: The connected sequence of prompts, outputs, edits, approvals, and retention actions that document the lifecycle of a communication or decision.

Proof Layer: The technical controls, logs, and records that enable organizations to validate, reconstruct, and defend AI-assisted activity in response to regulatory or legal inquiry.


Regulated Workflows: Business processes, communications, or decisions subject to regulatory requirements for supervision, retention, and auditability.

Continuity: An unbroken chain of evidence from initial AI involvement through to retention and production.

Research conducted by




Industry distribution



Region distribution




Function / department




Seniority / role level




Organization size



Annual revenue



Decision authority split



Research conducted by

About Arctera

Arctera, a Cloud Software Group company, is the leading global provider of compliance and governance solutions that enable firms to unleash game-changing technologies into their organizations while minimizing risk. Created in 2024 from Veritas Technologies, Arctera helps the biggest companies in the world monitor and control exactly how their information is being accessed, used and shared. The Arctera Unified Platform is able to capture data from over 130+ different content sources, and more than 280 AI policies help firms streamline compliance and adapt to evolving regulations.

Copyright © 2026 Arctera. All rights reserved. Arctera and the Arctera Logo are trademarks or registered trademarks of Arctera or its affiliates in the U.S. Other names may be trademarks of their respective owners.



Learn more at arctera.com

Connect with us on [LinkedIn](#)

Contact: press@arctera.com