zoom

WHITEPAPER

Navigating the compliance risks of generative AI in financial services

UCC Leaders hesitating to deploy AI can use Zoom Compliance Manager powered by Theta Lake to help mitigate risk

zoom
Compliance Manager

POWERED BY THETA LAKE

Introduction

The deployment of generative artificial intelligence (GenAI) across enterprise environments is accelerating, driven by demands for greater productivity, enhanced user experience, and intelligent support across commu- nication workflows. Nowhere is this acceleration more fraught than in financial services, where strict regulato- ry expectations intersect with rapidly evolving collaboration platforms and user expectations.

Unified Communications and Collaboration (UCC) leaders are under increasing pressure to modernize the digital workplace while being asked to maintain control over data, disclosures, and recordkeeping. This paper addresses that tension, especially as it applies to Zoom AI Companion, and examines how Zoom Compliance Manager (ZCM) powered by Theta Lake enables firms to adopt AI capabilities responsibly.

The AI boom and compliance concerns

As financial services firms explore new frontiers in operational efficiency, the adoption of GenAI has emerged as both a promising innovation and a source of considerable regulatory anxiety. Tools such as Zoom AI Companion offer the potential to streamline workflows through real-time transcription, summarization, and intelligent suggestions. These capabilities speak directly to the productivity goals of modern UCC environments and for firms seeking to "Work Happy" in an increasingly digital world, the appeal of these tools is clear.

However, the regulatory environment in which financial institutions operate is uniquely stringent. Every communication, whether spoken, written, or algorithmically generated, may be subject to supervision, recordkeeping, and retention requirements under frameworks such as SEC Rule 17a-4 and FINRA Rule 3110. In this context, the introduction of generative AI tools into daily operations is not simply a matter of user enablement. It is a question of governance.

The consequence is an industry-wide hesitancy. Despite growing awareness of the competitive advantage AI-enhanced communications can bring, compliance teams and Unified Communication and Collaboration (UCC) leaders remain cautious. As Theta Lake's Digital Communications Governance & Archiving Compliance & Security Report 2024/25 notes, 97% of financial firms see significant risks associated with AI use. Data privacy, misinformation, and regulatory exposure are among the leading concerns .

Why this paper matters

This paper examines the tension between innovation and obligation and between the productivity gains offered by generative AI and the compliance realities that constrain its deployment in financial services. It provides a structured analysis of the capabilities and risks of AI in UCC environments, outlines the regulatory mandates shaping AI adoption, and identifies the core reasons firms hesitate to embrace AI tools.

Most importantly, it highlights how Zoom Compliance Manager (ZCM) powered by Theta Lake, a compliance and risk platform natively integrated into the Zoom environment, offers a path forward. By aligning AI- enhanced communications with existing supervision, archiving, and data governance frameworks, ZCM helps mitigate key risks while enabling organizations to adopt AI features responsibly. In an era where communication is both an asset and a liability, such tools are not optional, they are essential.

97%

of financial firms see significant risks associated with AI use.

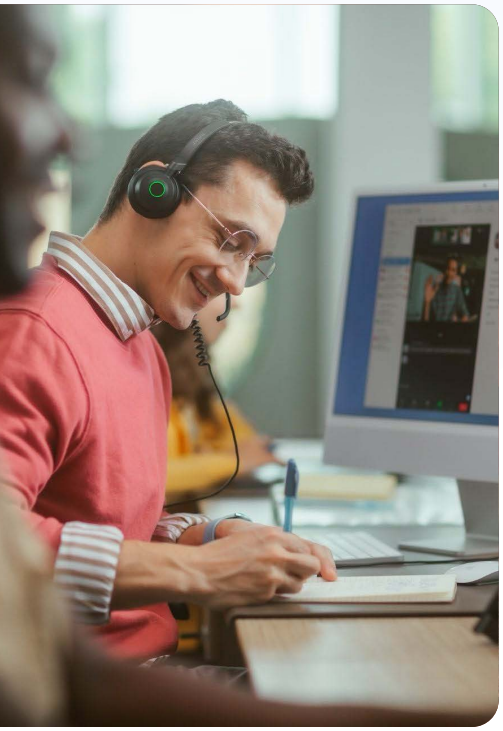

Generative AI in unified communications: capabilities and risks

The use of generative AI in collaboration tools represents a major shift in how financial services professionals communicate, document, and engage. What began as simple meeting transcriptions or text suggestions has expanded into full-featured AI assistance, capable of summarizing discussions, composing messages, and delivering contextual insights. These capabilities, while powerful, sit at the intersection of two competing demands: the drive to enable seamless, productive work, and the obligation to monitor, retain, and supervise business communications within a regulated framework. Understanding both the potential and the risks of these technologies is critical to navigating this rapidly evolving landscape.

The role of generative AI in UCC platforms

In recent years, UCC platforms have evolved from basic meeting tools into dynamic work hubs that facilitate hybrid engagement across voice, video, chat, file sharing, and productivity integrations. With that evolution, generative AI has emerged as a central feature that's embedded to support more efficient, contextual, and responsive collaboration.

Zoom AI Companion exemplifies this integration by offering features such as real-time meeting summaries, automated action-item generation, chat composition assistance, and the ability to "catch up" on missed threads with AI-curated recaps. These capabilities are not mere conveniences, they are powerful productivity multipliers. In highly regulated sectors like financial services, they hold the potential to assist with compliance- adjacent functions such as documentation, client follow-up, and communication transparency.

For regulated firms under pressure to do more with leaner resources, the promise of AI-enhanced UCC tools is compelling. However, the compliance risks they introduce are significant and cannot be overstated.

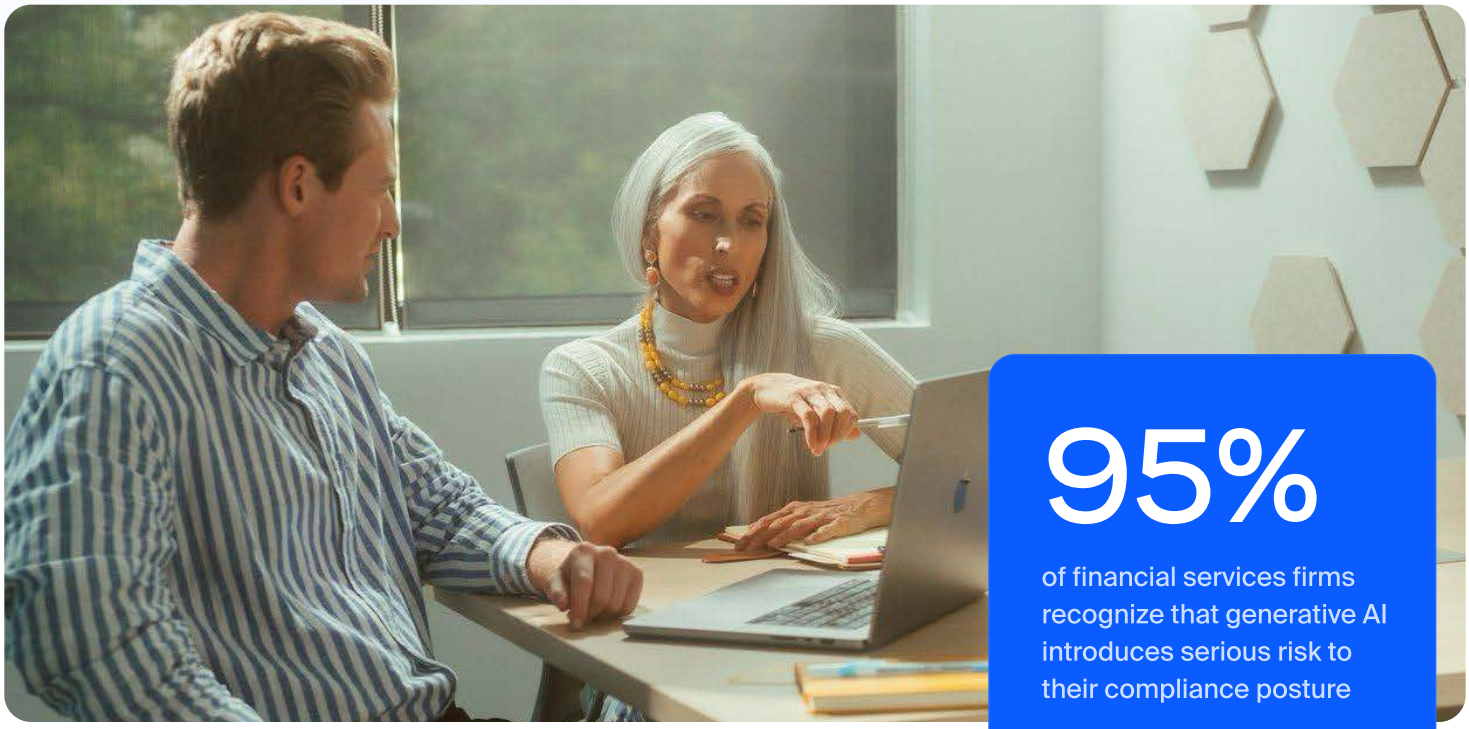

95%

of financial services firms recognize that generative AI introduces serious risk to their compliance posture

Key compliance risks identified by financial firms

According to Theta Lake's Digital Communications Governance & Archiving Compliance & Security Report 2024/25, 97% of financial services firms recognize that generative AI introduces serious risk to their compliance posture . Over half of all respondents (53%) cited the sharing of confidential data with external AI tools and large language models as the greatest concern. In effect, this reflects a fear of data escape, where AI engines might be trained on or exposed to sensitive, non-public information.

The nature of AI-generated content itself presents a second layer of risk. Thirty-eight percent of respondents expressed concern that meeting summaries generated by notetaker bots or other AI tools could lead to regulatory exposure.

In regulated environments, even automatically generated records must meet supervision, retention, and accessibility requirements. When firms are uncertain whether AI-generated content is being properly captured, stored, or governed, it raises the possibility of compliance gaps, regardless of the summaries' accuracy or utility.

Additional risks include AI-generated outbound communications, chatbot interactions, and the summarization of internal datasets; each of which raises concerns about misinformation, data integrity, and oversight. Importantly, even if AI outputs are accurate in form, they may still violate advertising standards or supervision requirements if presented without human review.

62%

of firms are already using machine learning or natural language processing (NLP) in their supervision processes

34%

of firms working to improve their underlying datasets

28%

of firms describing their implementations as resource-intensive.

Fragmented oversight: AI in compliance workflows

Compounding these risks is the challenge of applying AI within compliance monitoring itself. The same Theta Lake survey reveals that 62% of firms are already using machine learning or natural language processing (NLP) in their supervision processes. Yet, a majority of these firms report significant challenges due to fragmented data sources and incomplete record sets, with 34% working to improve their underlying datasets and 28% describing their implementations as resource-intensive.

This fragmentation not only limits the efficacy of AI in identifying risk, it undermines trust in the governance frameworks meant to contain that risk. Without unified, policy-enforced capture of all communications modalities (including AI-generated outputs), organizations risk falling out of compliance even as they attempt to modernize their oversight functions.

The regulatory and compliance barriers to AI adoption

The increasing availability of generative AI in unified communications platforms poses a direct challenge to the regulatory mandates that govern financial services institutions. While these tools promise operational efficiency, they do so within a legal environment that holds firms accountable for every business-related communication, regardless of whether it is authored by a person or an algorithm.

Key Compliance Challenges in AI Use for Financial Services

Among the most pressing concerns is the potential exposure of sensitive data. Generative AI tools frequently rely on access to user input, system content, or internal knowledge bases to function effectively. In the context of financial services, this raises questions about whether confidential client information could be inadvertently shared, processed, or retained in ways that breach privacy obligations.

According to Theta Lake's Digital Communications Governance & Archiving Compliance & Security Report 2024/25, 53% of respondents named the sharing of confidential information with external AI tools and large language models as one of the top perceived risks associated with generative AI.

Equally problematic is the treatment of AI-generated content under supervisory and recordkeeping obligations. Under FINRA Rule 3110 and SEC Rule 17a-4, financial firms must retain and supervise shared communications that pertain to business operations, customer interactions, and investment recommendations. If AI-generated summaries, chat completions, or meeting transcripts fall within the scope of these communications—and are shared—they must be retained in their entirety, produced upon request, and subject to firm-level oversight. This introduces significant compliance risk. AI tools, by design, are capable of generating autonomous or semi-autonomous content that may not be reviewed before being acted upon or distributed. A misinterpreted phrase, an omitted disclaimer, or an AI-generated summary that mischaracterizes a discussion can result in an inaccurate or misleading record. In such cases, the firm, not the algorithm, bears regulatory responsibility.

These concerns are compounded by a broader lack of regulatory clarity. Few formal guidelines exist on how to apply traditional supervisory expectations to generative AI outputs. As a result, compliance teams are forced to interpret existing frameworks within novel technological contexts, often defaulting to the most risk-averse option: disabling AI functionality altogether.

FINRA and SEC expectations for AI-generated communications

Despite the lack of AI-specific regulations, the expectations from U.S. regulators are unambiguous. FINRA has reiterated that firms are accountable for communications regardless of their origin. Whether an employee or an AI tool generates a statement, if it conveys a message related to a firm's business, it must be retained and supervised accordingly.

This includes adherence to the "fair and balanced" standard in FINRA Rule 2210, which prohibits exaggerated or misleading claims in retail communications. Any AI-generated summaries, promotional messages, or responses used in customer-facing contexts must meet the same standards as human-generated materials. Similarly, SEC Rule 17a-4(f) sets strict requirements for the technical treatment of electronic records, including mandates for non-rewriteable, non-erasable formats and reliable audit trails.

These requirements place a high burden on firms adopting AI-enhanced tools. Unless the AI outputs are fully captured, stored in compliant formats, and integrated into existing supervision workflows, their use can trigger violations, even if those outputs never leave the firm.

Why financial leaders are reluctant to deploy AI Companion

Despite growing recognition of the efficiency gains generative AI technology offers, a considerable number of financial institutions are choosing not to deploy them within their unified communications environments. Zoom AI Companion, with its user-friendly ability to generate meeting summaries, compose chat responses, and provide real-time contextual insights, presents precisely the type of innovation firms are eager to embrace, yet hesitant to trust. The hesitation, as revealed in Theta Lake's **Digital Communications Governance & Archiving Compliance & Security Report 2024/25**, is neither vague nor speculative. It is grounded in the risks and regulatory pressures that define financial services operations.

Survey insights: the AI compliance hesitation in financial services

According to the report, more than half of financial services firms (54%) have actively disabled AI features in their communications tools due to compliance, privacy, or search visibility concerns. These aren't fringe capabilities. Among the top features being restricted of financial services firms have actively disabled AI features in their communications tools due to concerns 54% are AI note-takers, app integrations, and message reaction functionalities such as emojis. This data suggests a broad institutional trend toward technological conservatism, driven not by a resistance to innovation, but by a lack of confidence in the regulatory defensibility of AI-driven outputs.

This hesitancy reflects a practical concern: generative AI outputs do not yet fit neatly within the supervision and retention models financial services firms rely upon. In environments where every client communication may be subject to audit or litigation, uncertainty over whether AI-generated content is properly retained, governed, and supervised creates a disproportionate compliance risk.

54%

of financial services firms have actively disabled AI features in their communications tools due to concerns

In an effort to avoid these risks, many organizations are sacrificing the very tools designed to improve engagement and productivity. Yet this approach creates new problems. Disabling features may lead to reduced employee productivity (reported by 42% of firms), impaired compliance team efficiency (46%), and an increase in off-channel communications (47%), each of which introduces its own set of risks and inefficiencies.

$4B

in fines levied against financial firms*

Potential legal and regulatory fallout

The consequences of insufficient communications governance are not hypothetical. Over the past three years, nearly $4 billion in fines have been levied against financial firms for recordkeeping and supervision failures. And it's a trend that shows no sign of reversal. As regulators continue to expand the scope of their enforcement, particularly in relation to modern communication modalities such as chat, video, and screen sharing, the inclusion of AI-generated content in regulatory inquiries is inevitable.

Firms that implement generative AI tools without a robust compliance strategy risk introducing exactly the kinds of vulnerabilities that have led to multi-million-dollar enforcement actions in recent years. While the productivity benefits of tools like Zoom AI Companion are real, so too are the legal liabilities of deploying them without the infrastructure to govern their use.

*Over the past three years

Uncertainty about AI decision-making

A final barrier to adoption lies in the opacity of AI decision-making itself. Generative AI models, while highly capable, do not yet possess the contextual awareness or regulatory fluency that human employees bring to communication. They are incapable of distinguishing between nuanced risk categories, interpreting firm- specific policy exceptions, or responding dynamically to evolving compliance guidance.

This epistemic gap and the absence of explainability in how AI arrives at its conclusions, renders many firms uneasy about allowing AI-generated communications to proceed without human intervention and forensic- level visibility into both inputs and outputs. In the absence of transparency, financial leaders are reluctant to accept accountability for content their compliance teams cannot fully audit or explain. For these reasons, the path forward for AI deployment must not only address feature functionality, but must also offer confidence in control, oversight, and defensibility.

How Zoom Compliance Manager powered by Theta Lake (ZCM) mitigates AI risks

ZCM is not merely a monitoring overlay or an external bolt-on. It is a compliance and risk governance platform architected to support the full lifecycle of communications within the Zoom ecosystem, including content generated by AI Companion. By embedding surveillance, supervision, and archiving capabilities directly into the platform's native workflows, ZCM helps ensure that AI adoption does not compromise regulatory posture.

What makes ZCM different for AI governance

Unlike conventional compliance platforms, ZCM is purpose-built to help customers supervise modern communication content using the same archiving, retention, and review processes applied to human-created communications.

With AI-powered risk detection and native integration with Zoom, ZCM gives compliance teams the ability to confidently govern their communications, regardless of origin. This is not a workaround for AI, it is compliance infrastructure designed for the way people work today.

Ensuring secure AI use within compliance boundaries

Financial services institutions require more than AI functionality. They require enforceable control. ZCM enables administrators to define policies that govern where and how AI features are enabled across users and meeting types. These controls empower firms to align AI usage with their risk tolerance, restrict access to sensitive workflows, and ensure that all communications (including those supported by AI) remain subject to enterprise-level governance.

Zoom AI Companion is designed with privacy and transparency principles in mind. For example, it does not use customer audio, video, chat, or other communications-like content to train Zoom's or its third-party AI models. However, the outputs of AI Companion must still be retained, discoverable, and supervised to meet compliance expectations. ZCM facilitates this by enabling firms to automatically capture and store AI- generated summaries so they can be archived according to SEC Rule 17a-4 and FINRA Rule 3110 requirements. Support for AI-generated responses is expected to follow.

ZCM also addresses a deeper, often unspoken concern among compliance leaders: that AI-generated content cannot be governed with the same level of rigor as traditional communications. This skepticism stems not only from a lack of regulatory clarity, but also from the limitations of generic compliance tools that were not designed to recognize or supervise AI outputs. ZCM, by contrast, applies compliance controls to AI-generated content using the same precision and policy logic applied to human communications, so supervision is not compromised, but extended.

Addressing core compliance concerns

The risk of non-compliance does not stem solely from the AI tool itself, but from the absence of visibility, auditability, and oversight.

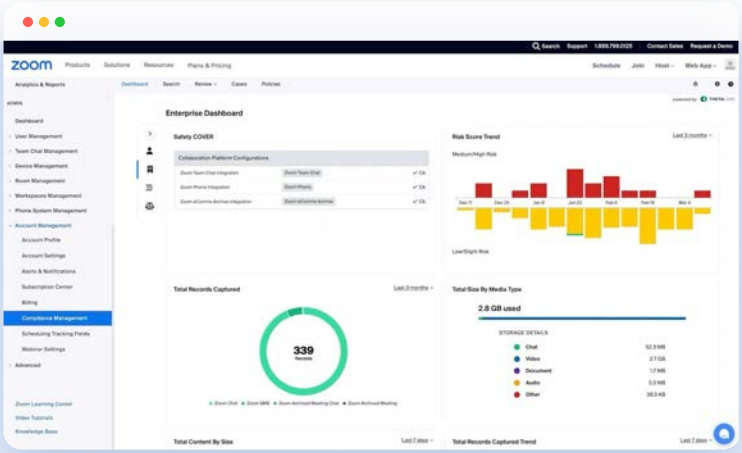

ZCM simplifies compliance review with AI-powered risk detection and forensic-level visibility into the content that AI Companion generates.

ZCM addresses these challenges in three primary ways:

01

Automated risk detection

ZCM applies pre-configured, policy-based classifiers to detect risky language, data exposure, or potential misconduct in communications, including AI-generated summaries and messages. These classifiers are designed to align with regulatory frameworks such as SEC and FINRA rules, and are continually refined and tested by Theta Lake to reflect evolving industry standards. Organizations can tailor classifier sets to their risk tolerance, business lines, or jurisdictions, and deploy them easily through a graphical interface, without the need for custom code. With support for over a dozen languages and AI-driven explainability features that help reviewers understand why content was flagged, ZCM empowers compliance teams to detect and respond to risk across formats and functions with precision.

02

Compliant archiving and secure retention

All captured content, including AI-enhanced summaries, is retained in formats that meet regulatory standards for immutability, accessibility, and searchability. ZCM leverages Theta Lake's SEC Rule 17a-4(f)- compliant archive ,to preserve records in non-rewriteable, non-erasable formats with time-stamped and serialized storage. Data is encrypted both in transit and at rest, with support for customer-managed encryption keys and strict access controls to prevent unauthorized retrieval. Theta Lake's infrastructure is certified under SOC 2 Type II, ISO 27001, and other leading security frameworks, supporting data integrity, privacy, and regulatory defensibility. ZCM also enables comprehensive eDiscovery and legal hold capabilities, so AI-generated content is included in search and supervision workflows alongside all other modalities.

03

Configuration controls and workflow integration

ZCM allows firms to enforce policy-based usage restrictions and integrates with supervisory review workflows to ensure that AI-generated content can be reviewed alongside other records. These capabilities help firms avoid creating shadow data or uncontrolled content sets that fall outside the purview of compliance. Together, these features enable financial services firms to engage with Zoom AI Companion not as a liability, but as a governed asset—one that can be deployed with confidence in even the most regulated environments.

The path forward: best practices for deploying AI in financial services

The adoption of generative AI in unified communications is not a question of if, but how—and when. For financial institutions, the path forward must be grounded in a strategy that reconciles innovation with the full scope of regulatory obligations. The objective is not to delay transformation, but to enable it responsibly: with the necessary policies, oversight, and infrastructure in place to safeguard the institution, its clients, and its data.

AI solutions like Zoom AI Companion introduce substantial opportunities to improve efficiency, documentation, and client engagement. Yet these advantages must be harnessed within a governance framework that ensures compliance is not compromised in the process. ZCM provides the operational foundation on which such a strategy can be built.

To move from caution to confident adoption, financial services firms should consider the following best practices:

01

Implementing AI with strong compliance controls

The deployment of generative AI should be accompanied by clear, documented policies outlining where, when, and how AI features may be used in communications workflows. These policies should distinguish between internal experimentation and client-facing use, and should incorporate supervisory review wherever AI-generated content may carry regulatory relevance.

ZCM enables organizations to enforce these policies through fine-grained configuration controls, allowing UCC leaders and compliance teams to manage AI feature availability in alignment with risk tolerance and regulatory guidance. In parallel, organizations must ensure that all AI-generated outputs, such as meeting summaries and action items, are captured, archived, and made searchable alongside other business records.

02

Future-proofing AI adoption for compliance readiness

Because both AI technology and regulatory expectations are evolving rapidly, financial firms must take a future-ready approach to communications governance. This means investing in infrastructure that supports consistent policy enforcement across modalities and platforms, and that enables quick adaptation when new communication features, or new regulatory interpretations emerge.

Central to this strategy is adopting a Digital Communications Governance and Archiving (DCGA) platform that integrates compliance oversight across chat, video, voice, file sharing, and AI-assisted content. Firms should ensure that AI governance is not siloed, but part of a unified compliance architecture that aligns with established frameworks for supervision, retention, and audit readiness.

ZCM's ability to extend compliance across all communication modalities, while also delivering AI-specific controls and risk detection, helps firms meet both current and emerging demands. It also ensures that communications compliance teams are equipped not only to monitor risk, but to respond to it—without obstructing innovation.

Conclusion

As generative AI becomes embedded in the day-to-day operations of financial services firms, leaders face a dual imperative: to accelerate innovation while safeguarding compliance. This paper has explored the regulatory barriers, operational hesitations, and governance solutions that define this moment.

What emerges is a clear path forward that does not require compromise, but rather coordination between technology enablement and regulatory rigor.

Balancing innovation with compliance

The financial services sector stands at a pivotal moment. Generative AI has become an integral part of the digital workplace, offering real-time insights, automation, and productivity enhancements through tools like Zoom AI Companion. Yet these capabilities arrive at a time of escalating regulatory expectations, mounting enforcement actions, and rising pressure on compliance teams to govern an ever-expanding volume and variety of communications.

The answer is not to delay AI adoption, but to pursue it with the same rigor and discipline applied to all other regulated activities. Compliance frameworks must evolve alongside communications technology. Zoom Compliance Manager (ZCM) powered by Theta Lake provides the control, visibility, and assurance firms need to responsibly enable AI across their Zoom environments. With native integration, SEC 17a-4(f)-compliant archiving, risk detection, and robust policy enforcement, ZCM helps financial services customers address the core barriers to adoption, allowing firms to deploy AI features without compromising regulatory posture.

Final takeaway for financial leaders

Firms can embrace the advantages of Zoom AI Companion without increasing compliance risk—if they adopt the appropriate governance infrastructure. By aligning AI deployment with established supervisory and retention mandates and leveraging purpose-built compliance solutions like ZCM, financial institutions can meet the moment with confidence, clarity, and control.