

March 31, 2025

The Secretary
Ontario Securities Commission
20 Queen Street West
22nd Floor, Box 55
Toronto, Ontario M5H 3S8

Me Philippe Lebel
Corporate Secretary and Executive Director, Legal Affairs
Autorité des marchés financiers
Place de la Cité, tour PwC
2640, boulevard Laurier, bureau 400
Québec (Québec) G1V 5C1

Submitted via email:
comments@osc.gov.on.ca
consultation-en-cours@lautorite.qc.ca

**Re: Theta Lake, Inc.'s Feedback on CSA Staff Notice and Consultation 11-348 -
Applicability of Canadian Securities Laws and the use of Artificial Intelligence Systems in
Capital Markets**

Dear CSA Members:

Theta Lake, Inc. ("Theta Lake") respectfully submits this letter in response to the request for comment ("RFC") on the applicability of Canadian securities laws and the use of artificial intelligence systems in capital markets.

In this letter, Theta Lake responds to Questions 4, 6, and 10 posed in the RFC and provides details of its AI-enabled platform as supporting context and examples. Our goal is that Theta Lake's response to the RFC helps the CSA to consider how AI is used for the purpose of compliance and security oversight to inform regulatory analysis. At present, we believe the CSA and CIRO regulatory frameworks for compliance and risk management are sufficient to address the risks of AI.

Theta Lake Overview

Theta Lake provides a multi-patented, AI-enabled, Digital Communications Governance and Archiving ("DCGA") platform that assists CSA-regulated entities to meet obligations for the capture, retention, search, and proactive compliance of electronic communications ("e-comms"). Theta Lake's DCGA platform supports compliance with rules such as CSA's National Instrument 31-103 Section 11.5(2)(n) requirement to "document correspondence with clients," MFDR Rule 5.1(o), IDPC Rule 3804(2)(xiv), and the related communication review and supervisory mandates of IDPC Rule 3602 and its equivalents.

Theta Lake has been named as the leading Visionary in the inaugural [2025 Gartner® Magic Quadrant™](#) for Digital Communications Governance and Archiving Solutions and ranked number one in the [Regulatory Compliance Critical Capabilities Use Case](#). Theta Lake currently supports recordkeeping and supervision compliance for CSA-regulated firms as well as those subject to equivalent USA, UK, EU, Australian, Hong Kong, and Singapore rules.

Theta Lake integrates with over 100 communications platforms including Bloomberg, Microsoft Teams, NICE, Verint, Webex, WhatsApp, Zoom, and more, to enable compliance and security controls to support the cloud native, device agnostic modern workforce. Theta Lake counts Cisco, RingCentral, Salesforce/Slack, Zoom, and senior executives from Microsoft as its strategic investors.

These direct investments and seamless product integrations demonstrate a meaningful validation of our approach to DCGA and set us apart in the market segment. Ultimately, Theta Lake's unified capture, search, and analysis capabilities allow firms to deploy their communications platforms of choice, for the benefit of employees and customers, by facilitating AI-driven compliance at scale. In fact, Theta Lake was recently selected for the UK Financial Conduct Authority's [AI Spotlight](#) which showcases examples of how AI can drive positive change in financial services.

Question 4: Human oversight of AI systems

Theta Lake considers having a “human-in-the-loop”, where users can effectively monitor the input and/or output of an AI system, as an essential component of the design and use of its platform. At a foundational level, Theta Lake places direct human involvement at the center of our AI-enabled processes.

Theta Lake's AI-powered detections identify myriad compliance, privacy and security risks in captured communications. So, for example, Theta Lake can ingest a Zoom meeting where the chat and screen share capabilities are used and apply its AI models to the conversation to detect risks like potential collusion in the chat, or the display of a prospectus or sensitive personal data in the screen share. Risks identified by the AI models are presented to human reviewers for confirmation and analysis. Theta Lake visually pinpoints where its AI models have discovered a risk by labeling or highlighting the relevant section of a conversation so that reviewers can quickly spot and analyze key portions of the interaction. This enables the human oversight role to be incorporated directly into a firm's AI governance approach, allowing a compliance team's expert judgment and context-specific understanding to validate AI-identified risks.

Theta Lake further promotes explainability and human in the loop design through its detection annotation feature. Theta Lake users can view a plain language explanation of the rationale for the triggering of a particular AI risk detection directly in the platform. Using the Zoom example above, a user can request an annotation of the portion of the chat conversation triggering the collusion detection. The annotation might state that the chat conversation included the phrase “keep this between you and me” and the zipper face emoji, which suggest that the participants are attempting to obscure the conversation and keep it secret indicating potential collusion.

Annotations provide opportunities for education of supervisory staff and increase the accessibility and utility of the platform.

The skills and expertise of Theta Lake’s data science team is extensive. The team is staffed by scientists and engineers with distinguished academic backgrounds and decades of collective experience building machine learning models for financial services and other regulated industries. The data science group continuously develops and improves our set of over 100+ machine learning models. This expertise is reinforced by the regulatory intelligence team who assist in the AI model development process by offering feedback on relevant regulations and enforcement activity as well the interpretation of existing and prospective rulesets.

Question 6: Transparency and explainability of AI systems

We agree with the CSA’s view in the Staff Notice “that AI systems with the highest degree of explainability that is feasible in relation to the type of AI system being used will help promote transparency and better assist market participants in meeting their obligations under securities law.”

Foundational to its transparency practices, Theta Lake provides customers with full control over how AI detections operate within its platform.

Customers choose which AI risk detections are enabled and can activate or disable them at any time. Customers also manage the risk sensitivity of detections, categorizing them as risky, informational, or validations, which result in higher, lower, or neutral aggregated conversation risk scores. Risk scores allow compliance teams to prioritize how conversations are reviewed—high risk scores can drive near-time reviews of key conversations, whereas low- or no-risk interactions can be sampled on a percentage basis for spot checking purposes. The fine-grained transparency controls in Theta Lake’s platform allow organizations to understand, explain and customize AI behavior to address highly tailored compliance requirements.

Additionally, Theta Lake provides reporting on AI transparency and performance metrics as well as mechanisms for providing feedback about potentially anomalous results. Theta Lake’s classifier audit report provides measurements of hit rates for a given risk detection over time. So, for example, customers can query the number of Microsoft Teams meetings triggering the detection for misleading or inappropriate promotions across a specific time range and group of users. Customers can also provide direct feedback if models generate false positives or negatives. This feedback is used to improve the performance of Theta Lake’s risk detections to incorporate new topics or jargon within a particular risk domain.

The features discussed in this response to Question 6 are additive to the annotation capabilities described above.

Question 10: Existing vulnerabilities and systemic risks

As a third-party service provider to enterprise financial services organizations, Theta Lake understands the primacy of robust, auditable security protocols, many of which directly support AI development. Theta Lake’s cybersecurity and audit practices broadly align with the expectations outlined in CSA Staff Notice 11-332 on Cyber Security as well as the third-party service provider oversight requirements described in Part 11 of the CSA’s National Instrument 31-103.

Theta Lake continuously assesses AI vulnerabilities and systemic risks as part of its annual SOC 2, Type II, PCI DSS, and TruSight audits. Theta Lake maps its SOC 2 security controls to those of ISO 27001 and HIPAA each year.

AI risk and platform vulnerabilities are assessed each quarter as part of routine testing and as part of an annual independent third-party penetration test. Theta Lake's Software Development Lifecycle Process applies to its AI features ensuring that new risk detections, updates to existing models, and the rollout of new AI features are appropriately assessed.

These third-party audits provide consistent, repeatable, measurable protocols that demonstrate compliance with articulated controls, including those critical to AI development processes. The audits offer a standardized metric with which to assess third parties for whom AI is a core service offering. Firms conducting vendor due diligence can leverage these audit reports as an essential component of risk assessment processes.

Theta Lake's AI is designed to be fair and objective in its risk detection capabilities to mitigate widespread systemic bias. The AI models are trained on thousands of positive and negative examples of the risks they aim to detect, without incorporating characteristics that may introduce bias, such as age, race, religion, sexual orientation, or ethnicity. This means that the AI-enabled detections are focused solely on the contextual meaning of a conversation and the potential risks involved. The risk detections are not developed for high-risk purposes like underwriting, determining capital and reserve adequacy, investment advice, or credit determinations.

Conclusion

We hope that this background about how Theta Lake uses AI to support record keeping and compliance as well as insights into the human-in-the-loop approach to governance; transparency and explainability in AI systems; and mitigating risks, assists the CSAs as they examine the broad adoption of this new technology. As mentioned earlier, we do not believe that AI-specific rules or regulations are necessary at this time. We will continue to monitor this area to determine what additional steps we might take to iterate and improve on our risk management protocols. Theta Lake is in regular dialogue with US, UK and international regulators and would welcome the opportunity to discuss these issues with the OSC, AMF and CSAs to provide further context and demonstrate the value of AI for these evolving use cases.

Respectfully submitted,

/s/ Marc Gilman

Marc Gilman

General Counsel and VP of Compliance

/s/ Stacey English

Stacey English

Director of Regulatory Intelligence