THETALAKE

How Theta Lake Uses Artificial Intelligence in its Security and Compliance Suite

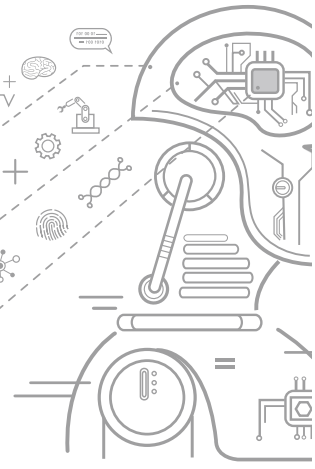

Theta Lake maintains 85+ artificial intelligence-based risk detections in its multi-patented Security and Compliance Suite, which are pre-trained and ready for customer use with no additional setup. Theta Lake's use of artificial intelligence ("AI") includes machine learning methods for natural language processing and computer vision to identify compliance issues like customer complaints, market abuse, and collusion across all communication modalities including video, voice, chat, email, and file transfers. Detections extend into privacy and cybersecurity risks such as identifying sensitive personal information (e.g., email addresses, birthdates, or account numbers), malware URLs, the display of sensitive applications like HR or finance tools, and the presence of financial logos, adult brands, hate speech, and offensive content.

Use of AI Models

There is no single "Theta Lake model." Theta Lake deploys AI models depending on the type of data being analyzed (video, voice, text, etc.) and the nature of the risk detection. Theta Lake combines traditional pattern matching techniques with the latest deep neural network methods to account for the weaknesses of legacy approaches to risk detection. For example, Theta Lake uses different models to detect a financial services logo in the background of a meeting video than it does to detect a chat conversation about off channel communications (e.g., "Let's take this offline on WhatsApp").

Customers Control AI Behavior

Customers have full control over how AI detections are enabled and operate within the Theta Lake platform. Controls include the ability to manage the risk sensitivity of detections to trigger as risky, informational, or validations, which contribute to a conversation's aggregate "risk score," and can be used to drive related supervisory processes. For example, customers can enable the detection for customer complaints to trigger as risky whenever conversations containing relevant language are identified.

As part of the human-in-the-loop process, conversations triggering risk detections are presented to reviewers for confirmation and further analysis. Using the example above, customers could configure a workflow requiring complaint conversations to be escalated to a specific compliance team for independent validation.

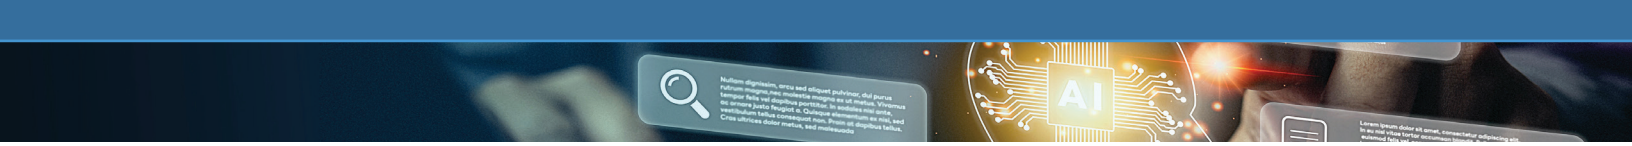

Customers can provide feedback on the results of Theta Lake's risk detections to further refine them and reduce future false positives and negatives.

THETA LAKE'S AI-ENABLED DETECTIONS DO MORE THAN SIMPLE KEYWORD MATCHING, AND IDENTIFY CONTEXT-SPECIFIC TERMS AND PHRASES.

Data Privacy, Exposure, and Loss Risks More information ›	Detection Media Types	Redaction Policy: None	Active	Risk
Background Risks				
Contracts and Contract Language -- Attachment and Emails			✓	
Credit Card and Driver's Licence Images -- Attachments			✓	
Digital Whiteboard Apps -- Videos beta			✓	
Notable Items Held up to Camera -- Videos			✓	
Office and Document Apps -- Videos			✓	
Online CRM & HR Apps and Portals -- Videos			✓	
Online Development & Infrastructure Apps and Portals -- Videos			✓	
Online Email Apps and Portals -- Videos			✓	
Online Financial Apps and Portals -- Videos			✓	
Whiteboards on Camera -- Videos beta			✓	

Model Training, Oversight, and Bias

A Theta Lake AI-based classifier is trained on thousands of positive and negative text and image examples of the type of risks it sets out to detect. As a very simple example, a classifier that detects conversations about off-channel communications is trained on positive texts like "let's take this call offline to WhatsApp", and negative texts, like "I'll be offline so can't accept your call." A machine learning model is trained to identify texts that are similar to the positive examples, and dissimilar to the negative examples.

Using AI avoids the pitfalls of lexicon-based solutions. Lexicons rely on a predefined set of rules or keywords to classify a given text, which lack the flexibility and adaptability needed to accurately classify texts in a complex communications environment. Theta Lake's AI-enabled detections do more than simple keyword matching, and identify context-specific terms and phrases.

For example, in the off-channel communications example above, Theta Lake's AI will detect the words "offlin" where the "e" is missing, "0ffline" where the first letter is the number zero, or "off mine" where the word has been improperly transcribed. Additionally, Theta Lake's AI will not trigger a detection for "offline" when it used to mean "unavailable" (e.g., "I'll be offline"), since it is context-sensitive. Using AI in this manner allows Theta Lake to identify and analyze imperfect or "noisy" data across documents, chat, email, video, and other visual content like images and emojis.

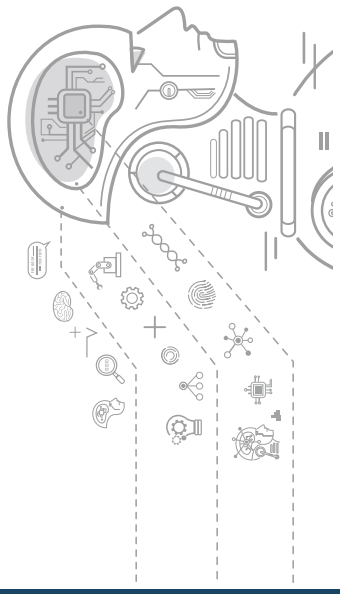

Theta Lake focuses on the long term performance of its models to minimize data and concept drift. Models are monitored post-deployment for unexpected changes and incorporate feedback from customers around false positives and false negatives. Models are refined and retrained routinely over time, incorporating new data sources in training data sets to account for emerging patterns. Theta Lake has several pending patents for model training and performance, including for optimal selection of training data, transcription analysis and normalization, fuzzy matching of words and phrases, and concept detection.

Theta Lake obtains training examples from various open, public datasets and does not use characteristics that may introduce bias into results such as age, race, religion, sexual orientation, or ethnicity when training its models. Theta Lake's risk detections are not used to make determinations about financial or behavioral activities such as underwriting, investment advice, human resources management (e.g., hiring, promotion, and termination), consumer lending, banking, or credit determinations. Moreover, Theta Lake does not attempt to infer sentiment or feeling, focusing only on the contextual meaning of a conversation and the potential risks.

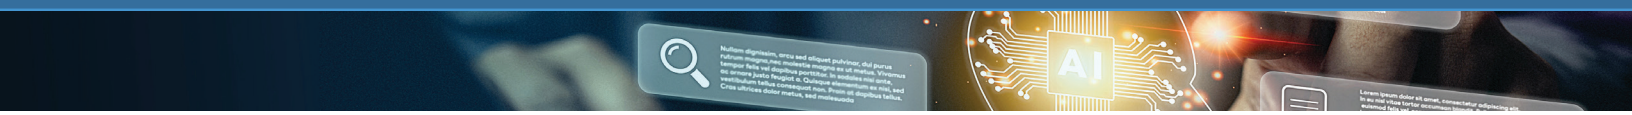

Compliance Evidence at Every Step

AN AUDIT TRAIL REPORT IS GENERATED FOR EVERY CONVERSATION INDICATING WHICH RISKS WERE DETECTED AS WELL AS PRESERVING ALL SUPERVISORY ACTIVITY FOR THE INTERACTION INCLUDING REVIEWER NAME, REVIEW DATE, COMMENTS, ESCALATIONS, ETC.

An audit trail report is generated for every conversation indicating which risks were detected as well as preserving all supervisory activity for the interaction including reviewer name, review date, comments, escalations, etc.. Audit trail reports provide compliance teams essential evidence of AI-enabled supervisory processes. Reports can be provided to regulators during examinations or investigations to demonstrate the rigor of the supervisory processes and the presence of a particular risk in a conversation.

Additionally, Theta Lake's built-in classifier audit report provides metrics around overall detection efficacy. For instance, customers can use the classifier audit report to query the number of Zoom video meetings triggering the detection for misleading promotions in the video, voice, chat, and file transfer communications across a specific time range or group of users.

The classifier audit report tracks the performance of risk detections over time and generates top level statistics regarding the frequency of positive or negative examples encountered across the organization, or within a particular communication platform. Classifier audit reports can be used as a feedback loop to track false negatives and positives. Reports can be provided to internal and external stakeholders to demonstrate AI oversight.

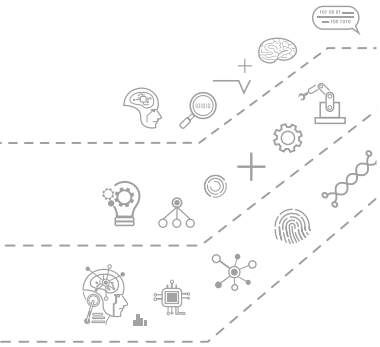

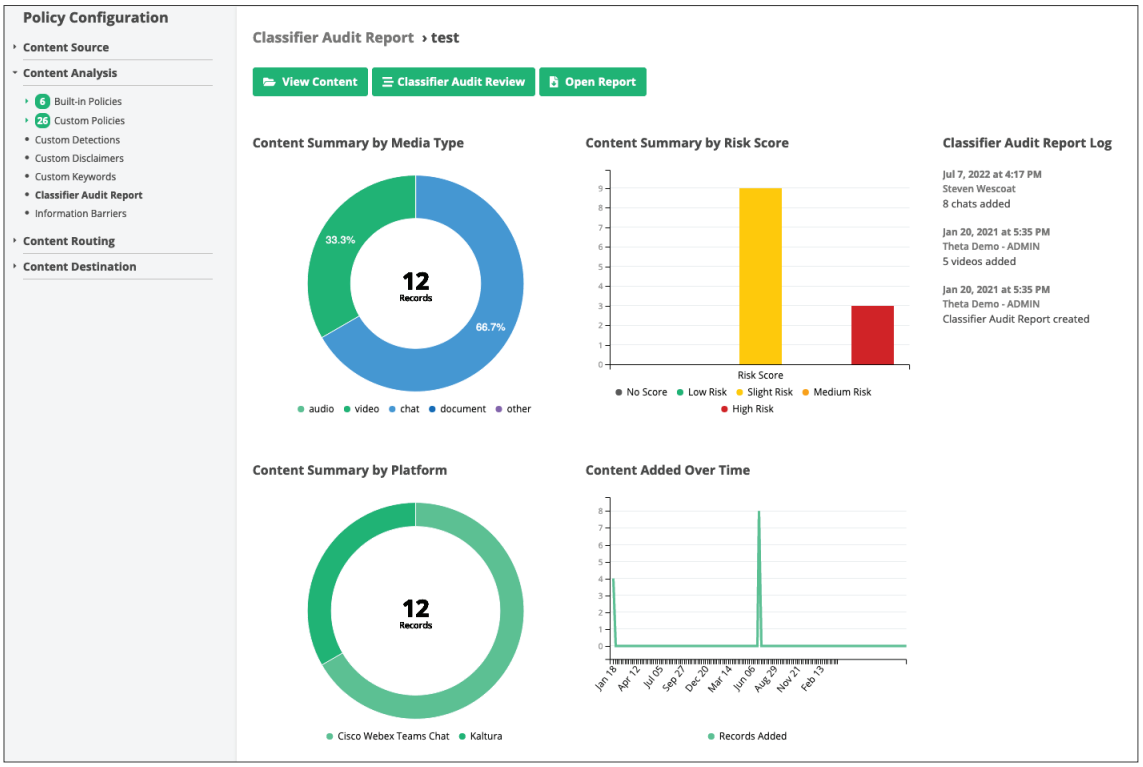

Learn more

To request a demo visit: https://thetalake.com/request-a-demo/

THETALAKE

ABOUT THETA LAKE. Theta Lake's multi-award winning product suite provides patented compliance and security for modern collaboration platforms, utilizing over 100 frictionless partner integrations that include RingCentral, Webex by Cisco, Microsoft Teams, Slack, Zoom, Movius and more. Theta Lake can capture, compliantly archive, and act as an archive connector for existing archives of record for video, voice, and chat collaboration systems. In addition to comprehensive capture and archiving, Theta Lake uses patented AI to detect and surface regulatory, privacy, and security risks in an AI assisted review workflow across what is shared, shown, spoken, and typed. Theta Lake enables organizations to safely, compliantly, and cost-effectively expand their use of communication platforms. Visit us at ThetaLake.com; LinkedIn; or Twitter at @thetalake.