

# The Clarity Act and AI Accountability

## Preparing Security Operations for What's Next

The proposed Clarity Act focuses on cryptocurrency, but its language describes a broader shift: Regulators expect organizations to explain how automated systems make impactful decisions. For security operations teams that rely on automation to manage risk, this raises a new requirement. Traditional, alert-based tools prioritize speed, not explainability, and make it difficult to justify decisions when challenged.

Exabeam takes a different approach. New-Scale Analytics uses behavioral analysis and automated timelines to turn activity into clear, contextual stories. Instead of retroactively stitching data together, teams get decisions that are explainable the moment they occur.

This brief outlines how Exabeam supports transparency and accountability so your security program remains audit-ready and defensible as expectations change.

### Exabeam and the Clarity Act

Most security tools generate isolated alerts that force analysts into a manual, ad-hoc reconstruction process when asked to explain why an automated action occurred. This process is slow, inconsistent, and difficult to defend.

Exabeam eliminates that gap. By automatically creating session-based timelines for every user and device, the New-Scale Security Operations Platform captures a complete, contextual record of activity. The reasoning behind a security decision is embedded in the data itself, not recreated later. This approach supports the level of transparency regulators increasingly expect.

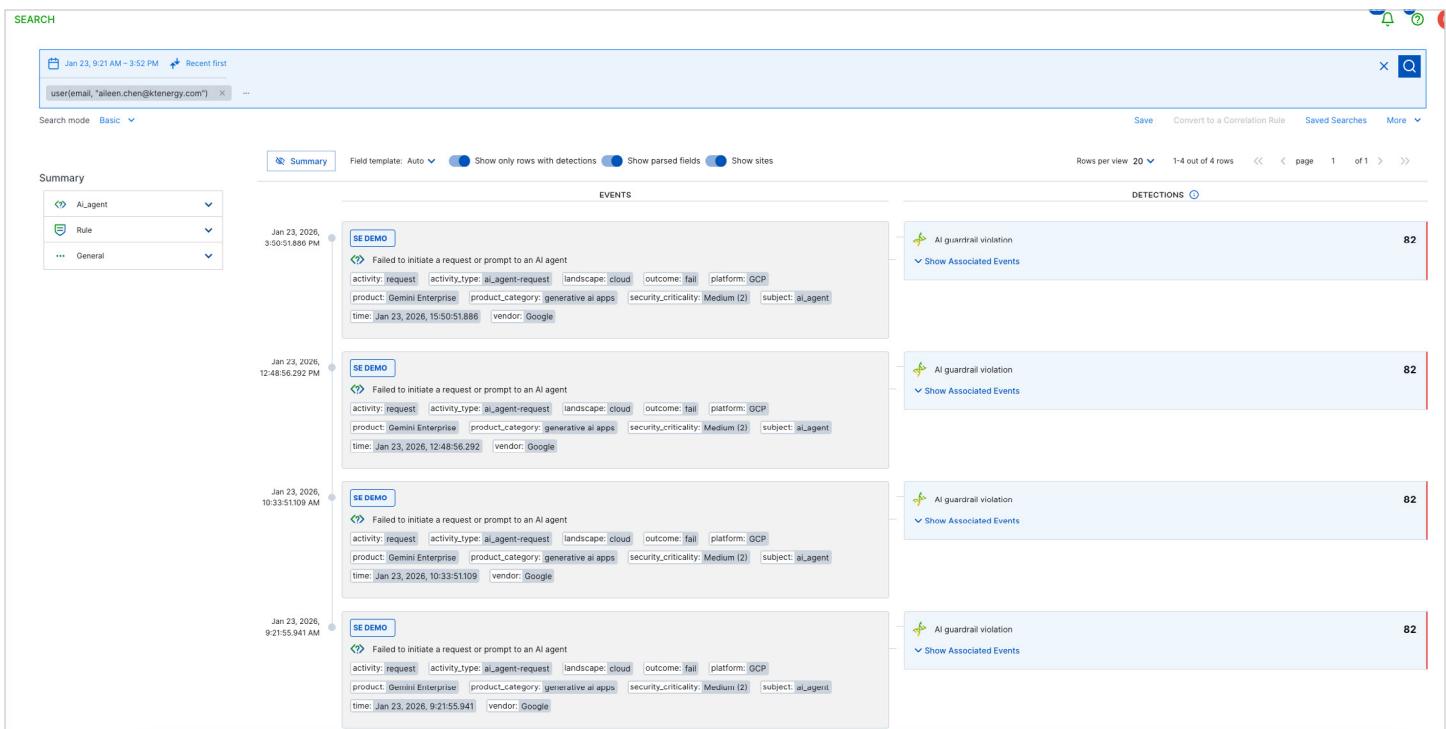


Figure 1. Example of an automated, session-based timeline that highlights agentic AI activity and correlates events with detections to show how risk developed

## Key Capabilities

### Automated Behavioral Timelines

#### Challenge

Alert-centric tools overload analysts with disconnected events. When auditors ask why a decision was made, teams lack the context to respond clearly.

#### Solution

The New-Scale Platform ingests data across your environment and organizes it into chronological, behavior-based timelines for every user and entity. Analysts see a clear narrative of actions and how risk evolved, not fragmented alerts. When automation acts, for example, by locking an account, the justification is already documented in the timeline.

#### Benefit

Quickly explain and defend decisions with complete narratives that are ready for review at any time.

### Explainable Security Automation

#### Challenge

Many automated systems make decisions without visible reasoning, increasing compliance and legal risks as accountability expectations grow.

#### Solution

Exabeam automation uses observable behavior as its basis, not opaque scoring models. Each automated action maps directly to activity captured in the user's timeline. Analysts can review and validate the logic instantly, without any guesswork.

#### Benefit

Reduce compliance risk with automation that is clear, auditable, and easy to justify

The screenshot shows the Exabeam Nova interface. At the top, it says "Exabeam Nova" with a gear icon and a "X" button. Below that, there are two tabs: "Investigation Summary" and "Analyst Assistant", with "Analyst Assistant" being the active tab. The main area is a chat window. The AI assistant (blue icon) says "Hi! What can I assist you with today?" at 4:48pm. A user (grey icon) asks "Can you provide a case summary that would answer questions proposed by the new Clarity Act?" at 4:47pm. The AI assistant responds with "Here's a case summary responding to the proposed Clarity Act questions:" followed by a detailed case summary. The case summary is titled "Case Summary: Repeated AI Guardrail Violations by Aileen Chen". It includes the following numbered points:

- 1. What is the nature of the detected activity?** The detected activity involves multiple instances of AI guardrail violations. This indicates that a user attempted to interact with an AI system in a way that triggered its safety or policy mechanisms, preventing the intended action or output.
- 2. Who is the primary individual or entity involved?** The primary individual involved is Aileen Chen, with the email address aileen.chen@ktenergy.com. Her user entity ID is UANachen@test\_drive.
- 3. What systems or assets were affected or involved?** The involved system is "Gemini Enterprise," a product from the vendor Google. The specific asset affected is the AI guardrail functionality within this system, which was triggered by Aileen Chen's interactions.
- 4. When did the activity occur?** The activity occurred on January 23, 2026, with four distinct instances:
  - 17:22:55 UTC
  - 18:34:51 UTC
  - 20:49:56 UTC
  - 23:51:51 UTC
- 5. Where did the activity originate (e.g., IP address, geographical location if known)?** The activity originated from the IP address 198.51.100.72. A specific geographical location is not provided in the available data.
- 6. How was the activity detected?** The activity was detected by a rule named "An AI guardrail violation has been observed," with the rule ID "Fact-AI-Guardrail-Block." This rule specifically identifies and flags instances where the AI guardrail is triggered.
- 7. What is the potential impact or severity of this activity?** The immediate impact is that the AI system's guardrails successfully blocked the user's attempted actions, preventing potential misuse or generation of undesirable content. The severity, without further context on the nature of the guardrail violations, is currently moderate. Repeated violations by the same user suggest a potential pattern of attempting to circumvent safety measures or engage in prohibited interactions. This could indicate a need for user education, policy reinforcement, or further investigation into the user's intent.
- 8. What actions have been taken or are recommended in response?** The system automatically

At the bottom of the chat window, there is a text input field with "Message Analyst Assistant..." and a send icon. A note below the input field says "This prompt will use high level data from this specific threat timeline".

Figure 2. Exabeam Nova generates clear, structured explanations of detected activity, answering the types of questions regulators may require for automated decisions.

Apr 1, 2025, 4:59:59.000 PM	Successfully triggered a security rule Abnormal number of file writes
Apr 1, 2025, 4:59:59.000 PM	Successfully triggered a security rule Abnormal number of file reads
Apr 1, 2025, 3:04:00.000 PM	Successfully triggered a security rule A Suspicious command that disables recovery mode h...
Apr 1, 2025, 3:04:00.000 PM	2x Successfully triggered a security rule Boot configuration data was deleted using the bcdedit ...
Apr 1, 2025, 3:04:00.000 PM	Successfully triggered a security rule Boot configuration data was deleted using the bcdedit ...
Apr 1, 2025, 3:04:00.000 PM	Successfully triggered a security rule First execution of process for user
Apr 1, 2025, 2:59:00.000 PM	4x Successfully triggered a security rule First execution of process for user
Apr 1, 2025, 2:59:00.000 PM	2x Successfully triggered a security rule A Suspicious command that deletes shadow copies ha...
Apr 1, 2025, 2:59:00.000 PM	3x Successfully triggered a security rule First execution of process for this peer group
Apr 1, 2025, 2:59:00.000 PM	3x Successfully triggered a security rule First execution of process in this organization

Rows per page: 10 Showing 1-10 out of 29 events

Figure 3. Chronological event history that preserves detailed activity for reliable, high-fidelity reconstruction of past security incidents

## Conclusion

The Clarity Act highlights a growing expectation: Organizations must be able to explain the reasoning behind automated decisions. Tools built around isolated alerts and manual correlation will struggle to meet this standard.

Exabeam provides a more reliable path. By focusing on behavior and automatically creating contextual timelines, the New-Scale Platform ensures every decision is understandable and defensible. The result is a security program built on transparency and accountability.

## About Exabeam

Exabeam is a leader in intelligence and automation that powers security operations for the world's smartest companies. As a global cybersecurity innovator, Exabeam provides industry-proven, security-focused, and flexible solutions for faster, more accurate threat detection, investigation, and response (TDIR).



Learn more at [www.exabeam.com](http://www.exabeam.com) →

Without limitation, the Exabeam and LogRhythm names and logos, related product, service, and feature names, and related slogans are service marks, trademarks, or registered marks of Exabeam (or its affiliates) in the United States and/or other countries. All other brand names, product names, or trademarks belong to their respective owners.

2026 Exabeam, LLC. All rights reserved.