212° DATA BOILER TECHNOLOGIES, LLC

BIG DATA | BIG PICTURE | BIG OPPORTUNITIES

We see big to continuously boil down the essential improvements until you achieve sustainable growth!

✉ info@databoiler.com https://www.databoiler.com

June 22, 2026

- via Electronic Mail (rule-comments@sec.gov)

Ms. Vanessa Countryman,
Secretary U.S. Securities and Exchange Commission (SEC)
100 F Street NE., Washington, DC 20549

Re: Concept Release on Consolidated Audit Trail and Other Audit Trails and Data Sources (File No. S7-2026-12; Release No. 34-105251; RIN 3235-AN54)[1]

Dear Ms. Countryman,

On behalf of Data Boiler Technologies, I am pleased to provide the SEC with our comments on the captioned release on Consolidated Audit Trail (CAT) and Other Audit Trails and Data Sources, in an attempt to inform any potential future agency action. Data Boiler is a Pioneer in FinTech with patented inventions in signal processing, trade analytics, machine learning, time-lock cryptography, etc. We frequently comment on regulatory policy both domestically and abroad with over 12 years in business. Please see the BOX below for an Executive Summary of our comments:

Opaque, Weaponized, Broken Beyond Repair and is wholly Inadequate (Annex 1)

(a) **Functional Mission Creep:** Congress's original mandate authorizing the CAT was explicitly limited to *flash crash prevention*; tellingly, there is no mention of *flash crash* anywhere throughout the SEC's entire concept release. The system has mutated from an emergency, post-2010 *flash-crash* mitigation utility, into an invasive, retroactive, permanent census – a sweeping extension of power that Congress never conferred upon the SEC.

(b) **Outdated Design:** With the CAT in its current centralized form, the database is highly vulnerable to security threats, unconstitutionally intrudes on everyone's privacy, and severely impair civil liberties of Americans who *transact* or *engage* in any way, shape, or form in the U.S. securities markets. The CAT architecture is fundamentally obsolete and structurally incapable of meeting twenty-first-century market realities. Framing its replacement as a choice between developing a different audit trail or relying on legacy data sources is a false dichotomy.

(c) **Fatal Flaws and Bias:** A so-called *"golden source single source of truth"* is indeed filled with *noises*, including initial bias, latency tolerance, human-induced opaque processes, institutional favoritism, poor controls, and practices that conceal or alter the essential order and trade sequences and artificial market events. The SEC, GAO, and DOJ must investigate if CAT data was *weaponized for political or commercial reasons* to hold the SROs accountable.

(d) **Undue Burden on Broker-Dealers + Misalignment of SEC Resources:** Total private-sector compliance expenditures to report CAT data now exceeds $1.7 billion annually. Continuous micro-technical rule modifications force broker-dealers to repeatedly waste millions of dollars on updating related systems. Dual-sided reporting can never be cost justified; it creates excessive data-in-motion traffic that is a costly waste and more susceptible to defects. Instead of being served by the CAT, massive data-cleansing and formatting inefficiencies trap highly paid agency personnel into serving as data processors administratively, rather than proactive gatekeepers of market integrity.

(e) **Flawed NMS Governance and Existential Structural Crisis:** Using an NMS plan to build a CAT creates severe conflicts of interest, in allowing for-profit SROs to diffuse operational responsibility. It is questionable for the CAT to continuously run basic data-wrapping loops on standard cloud servers to re-ingest public SIP and OPRA market data feeds – information that the SROs already natively generate and possess internally. The resulting regulatory cost-bloat essentially imposes an illegal Financial Transaction Tax on everyone, despite the Court vacating the funding order.

[1] https://www.sec.gov/files/rules/concept/2026/34-105251.pdf

Ultimately, no one wants to foot the bill for works they do not own, or to finance the hidden perks/ billable hours of opaque vendors, consultants, and lawyers contracted by the SROs behind closed doors.

(f) **"Everything Everywhere All at Once" Harms Everyone:** Modifying the *representative order linkage* requirements is a tacit admission that the original *daisy chain* approach was a flawed, incredibly expensive dead end. CAIS is an expensive and intrusive experiment. CAT was given 10+ years as an experiment, not once (Thesys), but twice (FINRA CAT LLC), to do a $2+ billion *proof-of-concept* that is doomed to failure. *Frequent pestering* (unlimited desires) to amend the CAT NMS Plan for an unrepairable CAT – or using stall tactics disguised as further reviews/ tests – erodes public trust.

(g) **Design and Scope Must Change:** Shift away from centralized data collection; adopt a Federated approach to fabricate the intelligent analytical layer. Have dedicated focus on volatility-event forensics and market-access risk controls. Expand product scope to futures, swaps, clearing data, and select digital-asset instruments. Require SEC and CFTC regulated SROs to supply the fastest, full-depth proprietary feeds. Exclude RFQs and primary-market activity. Shift lifecycle analysis to clearing & settlement systems, purging the entirety of CAIS. The CAT's unfixable structural flaws – including reliance on perimeter security, absent element-level protections, and vulnerable central administration – render it entirely incapable of meeting the Federal Zero Trust mandates; thus, it must be replaced (see Section H).

(h) **Out of Proportion, Revenue-Expense Mismatch, & Weaponization:** It is sad that human floor agents are being trusted less than AMM algorithms to maintain the continuous orderly function of markets. Port-level settings are important, but SROs using it as anti-competitive lock-in should be discouraged. SROs enjoys lucrative co-location profits while dodging responsibility to build a native compliance interface, and this thus represents an inappropriate cross-subsidization from CAT. LTID creates major enforcement risk – CAT reconstructs sequences using SIP & 3rd party data that inevitably drifts, producing false signals. It is unjust to shift the burden of proof to broker-dealers. This asymmetry enables weaponization of CAT for political or commercial reasons, undermines market integrity, and chill participation.

(i) **Do NOT Choose, Seek Alternative**: Hardening OATS/ COATS/ EBS – CapEx $250M-$450M depreciable over 7-10 years. CAT security patch $80M-$120M upfront CapEx + at least $40+M annually (**Caveat:** NOT avert cyber-honeypot risk). The only viable path is to abandon continuous centralized reporting entirely and shift to a federated model where data stays at its native source + deploy Agentic AI, restore need-to-know safeguards, and modernize market-monitoring.

(j) **Statutory Overreach, Bypassed Rulemaking Steps:** The SEC has mischaracterized the legacy EBS system to downplay CAT's far greater privacy and civil-liberties risks. Respect the EBS as a purpose-built insider-trading investigation tool with proper guardrails to ensure reasonable suspicion is established before summoning private data.

(k) **Unauthorized CENSUS, Civil Liberties and Privacy Violations:** Laundered massive government surveillance through SROs (contrasted to Vice President JD Vance's remarks about censorship), attempted to shield SROs from liability, and built a centralized repository that poses catastrophic national-security risks ($100M insurance cap grossly undermines a National security threats – a breach is not a minor corporate loss; it could trigger a structural collapse of U.S. capital markets). By enabling mission creep, policy circumvention, and the §31 fee extraction without proper rulemaking, the SEC and SROs effectively merged legislative, enforcement, and tax-collection powers – amassing authority beyond that of the U.S. President and undermining the separation of powers.

(l) Failed ZTA mandate, deliberately not following sound advice: Far riskier than building privacy-by-design. CAT SWG still relies on outdated NIST SP 800-53 Rev. 4 – "over a decade old" – leaving CAT vulnerable to modern cloud-native exfiltration and AI-driven reconstruction attacks. SEC/ SROs disregarded repeated warnings of a false sense of security, and the CAT became a fragile, bureaucratic honeypot that fails to protect markets or the public.

✉ info@databoiler.com https://www.databoiler.com

(m) **The CAT is Unsustainable:** Adding transparency measures or altering reporting formats will NOT resolve its underlying structural and constitutional vulnerabilities. We recommend shutting the CAT off immediately, patching OATS/COATS/EBS in the short-term, and going back to the drawing board, see Annex 2.

The Agentic Distributed Alternative (Annex 2) – See Blueprint drawing

(n) **Core Goals of a Modern Audit Trail Replacement:** The proposed architecture seeks to minimize the audit-trail footprint, reduce computational load, eliminate unnecessary trade reporting, and shift regulators from manual data processors to strategic gatekeepers. It emphasizes context-aware AI, Zero-Trust security, and selective ingestion of only high-value, anomaly-linked data. The system is designed to accelerate anomaly detection, reduce false positives, and identify emerging liquidity stresses or flash-crash precursors in real time.

(o) **Tier 1: Distributed, edge-based surveillance using AI Agents:** Tier 1 deploys localized AI agents at each SRO to run *"dual-track shadow processing"* that independently verifies exchange surveillance outputs. Raw matching-engine logs are analyzed in parallel by both the SRO's native tools and an independent AI agent, enabling immediate peer review and early anomaly detection. Manipulation patterns are decomposed into granular "triggers," allowing rapid matching against a machine-learning library, and drastically reducing computing resources while improving detection accuracy.

(p) **Tier 2: Semantic Audit Hub, Case Library, & QA Engine:** Tier 2 uses **Progress MarkLogic** and **Semaphore** to harmonize structured and unstructured data, resolve identities, perform semantic inference, and classify SRO misses into true positives, false positives, and false negatives. Only anomalies flagged by Tier 1 are unpacked, enriched with filings (e.g. 13F/13H), and converted into RDF triples for deep semantic analysis. This layer supports on-demand retrieval, automated mismatch resolution, and Zero-Trust access control, ensuring that raw data remains siloed at the source while enabling cross-market contextual understanding.

(q) **Tier 3: TPU-based reinforcement learning for cross-market stress detection:** Tier 3 escalates enriched anomaly packages to **Google TPU** clusters for reinforcement learning, cross-market correlation checks, and systemic-risk modeling. The TPU evaluates inventory imbalances, liquidity withdrawals, and multi-venue stress signatures to detect flash-crash precursors and refine SRO surveillance / volatility interruption mechanism parameters. Operating under strict Zero-Trust gating, it returns only cryptographic proofs and metadata to regulators, enabling real-time visibility without centralized data hoarding.

(r) **Advantages of the new architecture over CAT's centralized vault:** The three-tier distributed design eliminates CAT's honeypot vulnerabilities, reduces alert fatigue, and matches the pace of modern, cross-market manipulation by isolating identities and patterns across venues. It shifts surveillance from hindsight to active prevention, enabling dynamic recalibration of guardrails and volatility controls. By storing only anomaly-linked data and offloading heavy computation to targeted TPU jobs, the architecture cuts CAT's cloud-hosting costs by an estimated 70–80% while delivering far stronger security, privacy, and analytical capability than any centralized repository could achieve.

(s) ***A good decision, made now and pursued aggressively, is superior to a perfect decision made too late.*** Humans are slow and CANNOT manually reconcile the massive volume of structured trade logs and unstructured data driving modern markets. AI bridges this gap by handling tedious data ingestion and synthesis. Supported by human context, institutional knowledge, and strict guardrails, AI acts as a force multiplier – not job replacement. This shift elevates agency personnel from manual data processors to strategic gatekeepers of market integrity.

Disclaimer: Nothing contained in these submitted comments shall be construed as providing legal advice or establishing an attorney-client relationship, as we are not attorneys and do not offer legal counsel in any way, shape, or form. These

comments are submitted in good faith as a matter of public interest, exercising protected rights of speech and regulatory petition; under no circumstances shall this submission expose the author to civil liability, defamation claims, or retaliatory legal action by any party who may disagree with or object to the critiques expressed herein. Furthermore, under no circumstances shall this submission be interpreted as permitting the SEC, the SROs, or any other party to shift the operational, architectural, or intellectual burden of securing and organizing this multi-billion-dollar database away from the Commission and onto the public taxpayers and market participants who are forced to fund the CAT project.

Please see Annex 1 for our response to specific questions. Annex 2 is our recommended architectural design to replace CAT. Feel free to contact us with any questions and please keep us posted where our expertise might be helpful.

Sincerely,

Kelvin To

Founder and President

Data Boiler Technologies, LLC

This letter is also available at: https://www.DataBoiler.com/index_htm_files/DataBoiler%20SEC%20CAT%2020260622.pdf

CC:

Securities and Exchange Commission (SEC)
- The Honorable Paul S. Atkins, Chairman
- The Honorable Hester M. Peirce, Commissioner
- The Honorable Mark T. Uyeda, Commissioner
- Mr. Jamie Selway, Director of the Division of Trading and Markets
- Mr. David Hsu, Assistant Director, Office of Market Supervision, Division of Trading and Markets

Government Accountability Office (GAO)
- The Honorable Gene L. Dodaro, Comptroller General
- Mr. Michael E. Clements, Director - Financial Markets and Community Investment
- Dr. LaFountain Courtney, Director - Financial Markets and Community Investment

Department of Justice (DOJ) – Antitrust Division
- Mr. Omeed A. Assefi, Acting Assistant Attorney General
- Mr. David B. Lawrence, Policy Director

Commodity Futures Trading Commission (CFTC)
- The Honorable Michael S. Selig, Chairman
- Mr. Rahul Varma, Director of Market Oversight
- Mr. Jorge Herrada, Director of the Office of Technology Innovation

U.S. Department of the Treasury
- Mr. Gene Lange, Acting Under Secretary for Terrorism and Financial Intelligence
- Mr. Jonathan H. Burke, Assistant Secretary for Terrorist Financing and Financial Crimes
- Mr. Scott Rembrandt, Deputy Assistant Secretary for Strategic Policy, Terrorist Financing and Financial Crimes
- Ms. Anu Murgai, Director, Office of Capital Markets

Banking Regulators
- Dr. Rochelle M. Edge, Senior Advisor, Division of Financial Stability, FRB
- Ms. Amanda Freedle, Deputy Comptroller – Capital, Market Risk and Asset Management, OCC
- Mr. Ryan Billingsley, Director of the Division of Risk Management Supervision, FDIC

Executive Office of the President
- The Honorable Sean Cairncross, National Cyber Director

Table of Contents

ANNEX 1 – Data Boiler's response to specific SEC questions 8

Section A — Regulatory Purpose of the CAT 8

1. **What are the regulatory use cases that must be enabled for the Commission and the SROs to fulfill their statutory obligations?** Is the CAT necessary to enable those use cases? Are other audit trails or related data sources sufficient? Why or why not? 8

2. **Are there features of the CAT that could be eliminated** because they are unnecessary or could be replaced by other existing audit trails or data sources? Identify such features and alternatives. 10

3. **Should the Commission eliminate the CAT** in favor of developing a different audit trail or data source? Why or why not? How should a new system differ? What improvements could be gained? 12

Section B — Structure and Governance of the CAT 14

4. **What are the advantages and disadvantages of structuring the CAT as an NMS plan?** Should it continue to be structured this way? 14

5. **If not an NMS plan, how should the Commission and/or SROs direct and oversee the CAT?** What benefits, actions, and costs would be associated with transitioning? Would benefits offset costs? 15

6. **Should the Commission amend the CAT NMS Plan to implement a different voting structure?** What should it be and why? Should affiliated SRO groups or non-affiliated SROs receive extra votes? 16

7. **Should the Commission amend the CAT NMS Plan to require a different vote threshold** (majority, supermajority, unanimous) for any specific actions? Which actions and why? 17

8. **Are there measures that could increase transparency and accountability** around CAT NMS Plan voting while protecting sensitive information? Should more information be made public? 17

Advisory Committee 17

9. **What powers, responsibilities, or rights should be given to the CAT Advisory Committee?** How would these affect SRO governance? Would they raise conflicts of interest? 17

10. **Should membership of the Advisory Committee be reserved for certain interests or individuals with specific experience?** Should it be expanded to include ATS operators, technology experts, or others? How should representation be apportioned? 18

11. **Should the Advisory Committee have additional rights** beyond attending meetings and providing views? If so, what rights and why? 18

Section C — CAT Funding and Cost Management 19

12. **Should the Commission amend the CAT NMS Plan to require additional cost-management measures?** What measures and why? 19

13. **Should the Commission require the SROs to adopt a different funding model** for the CAT? What model and why? 20

14. **Should the Commission require the SROs to maintain reserve funds** for CAT operations? What size and purpose should such reserves serve? 21

15. **Should Section 31 fees or other alternative funding methods be used** to fund the CAT? What are the advantages and disadvantages? 22

Section D — CAT Design and Scope....23

16. **Should the Commission amend the CAT NMS Plan to modify the scope of CAT data collection?** What modifications and why?....23

a. Covered Securities (Product Scope) 23

b. Mandatory Reporting Entities (Jurisdictional Scope) 24

c. Covered Events (Lifecycle Scope) 25

d. Mandatory Data Points (Granularity Scope) 26

e. Explicit Exclusions (Out of Scope)....26

17. **Should the Commission modify CAT functionality** (e.g., lifecycle linkage, processing timelines, data retention)? What changes and why?....26

18. **Should the Commission modify CCID generation** or the way customer-level analysis is performed? What changes and why? 27

Section E — Previous Changes to CAT Requirements 28

19. **Should the Commission modify requirements related to verbal activity on exchange floors?** What changes and why? 28

20. **Should the Commission modify requirements related to electronic requests for quotes?** What changes and why?.. 29

21. **Should the Commission modify requirements related to port-level settings?** What changes and why? 29

22. **Should the Commission modify requirements related to representative order linkage?** What changes and why? 32

Section F — Potential Changes to Other Data Sources 35

23. **Should the Commission retire partially duplicative systems** (e.g., OATS, COATS, EBS)? Which ones and why? 35

24. **Should the Commission modify or replace the EBS system?** What changes and why? 37

25. **Should the Commission modify LTID requirements** or the large trader reporting system? What changes and why? 38

Section G — Civil Liberties and Privacy Considerations 39

26. **Are there additional privacy or civil liberties concerns** related to CAT or other audit trails? What changes should be made to address them?....39

27. **Should the Commission modify data-access controls** or other privacy protections? What changes and why? 42

Section H — Cybersecurity 45

28. **Should the Commission modify cybersecurity requirements** for CAT, EBS, LOPR, or other audit trails? What changes and why? 45

29. **Should the Commission require additional cybersecurity testing, audits, or certifications?** What requirements and why? 47

Section I — Transparency and Process of Comprehensive Review 48

30. **Should the Commission adopt additional transparency measures** regarding CAT operations, costs, or performance? What measures and why? 48

31. **Should the Commission modify the process for conducting comprehensive reviews** of audit trails and data sources? What changes and why? 49

Section J — General Request for Comment 50

32. **Are there any other considerations, costs, burdens, or benefits** related to audit trails or data sources that the Commission should evaluate? 50

33. **Are there potential regulatory responses not identified** in the release that the Commission should consider? 50

ANNEX 2 – Data Boiler's Vision of Next-Gen Market Monitoring System to replace CAT 52

Goals 53

The Meta-Regulatory Agent Topology 53

TIER 1 Intra-SRO Surveillance (**Peers Review**) Agents 53

TIER 2 Audit Trail Hub, Case Library, Quality Assurance (**Semantic Context Awareness**) 56

TIER 3 Reinforcement Learning Loops (**Cross-Market Validation**) 58

Advantages of the NEW 59

ANNEX 1 – Data Boiler's response to specific SEC questions

Section A — Regulatory Purpose of the CAT

1. **What are the regulatory use cases that must be enabled for the Commission and the SROs to fulfill their statutory obligations?** Is the CAT necessary to enable those use cases? Are other audit trails or related data sources sufficient? Why or why not?

   The last administration of the SEC relied extensively on CAT data to articulate (justify and quantify) proposed problems in its sweeping December 2022 Equity Market Structure Rule Proposals, which included:

   (a) Order Competition Rule: Proposed to require certain retail orders to be exposed to competition in qualified auctions.

   (b) Regulation Best Execution: Proposed to establish a heightened standard for broker-dealer best execution.

   (c) Disclosure of Order Execution Information (Rule 605 Update): Adopted in March 2024, the SEC used CAT data to analyze and expand the scope of execution quality reports for broker-dealers.

   (d) Minimum Pricing Increments & Access Fee Caps: Adopted to modify tick sizes and access fees, the SEC used CAT data to define the impacts on displayed liquidity.

   Reference to page 16 of our March 2023 comment letter, "the four proposal packages use four inconsistent rates for Attorney and Compliance Manager (see table on the right); … Practically, can every broker-dealer have dynamic price shopping capabilities like 'Booking.com' or 'Ticketmaster' – it is economically not viable. Even if given the sophisticated algorithms to determine 'where' and 'when' in indicating 'most favorable price', it is humanly impossible for the SEC examiners to objectively decipher if there may or may not be BestEx violation or conflicted transactions … Attempting to try use CAT for enforcement is like finding a needle in a haystack."[2]

   The SEC completely halted advancement on the two most controversial proposals (a) and (b) among the above 4 proposals following near-universal industry opposition. The market consensus and SIFMA comment feedback[3] suggested that the SEC must first collect and analyze data from the newly updated Rule 605 before it could ever legally justify bringing the Order Competition or Best Ex proposals back to life. For (c), the CAT data cited in SEC's proposal may simply be for the sake of the Administrative Procedure Act. Modernizing Rule 605 is obvious given it has not been updated since 2000. Yet, the effect of modernizing disclosure is unknown until the industry understand what changes will be in the National Best Bid and Offer (NBBO),[4] tick size (minimum pricing increment), and point (d) amended access fee caps (Rules 600(b)(89)(i)(F), 612, and 610(c)) are deferred until the first business day of November 2026,[5] and other possible change to Reg. NMS such as Rule 611 trade-through ban a.k.a. order protection),[6] etc.

   Per our Dec 26, 2025, comment letter,[7] "§17a-1 record retention requirements are obligations of the SROs that are separate from the private entity – CAT LLC (a subsidiary of FINRA). We oppose CAT LLC attempts to cross-subsidize SROs in fulfillment of obligations that deviate from the CAT project's original purposes. This could be considered functional creep or alleged misallocation of CAT funding resources."

[2] https://www.databoiler.com/index_htm_files/DataBoiler%20SEC%20Market%20Structure%20202303.pdf
[3] https://www.sifma.org/news/press-releases/sifma-comments-on-the-secs-equity-market-structure-proposals
[4] https://www.linkedin.com/pulse/us-market-data-reforms-complicated-cumbersome-comprehensive-kelvin-to-lzbse/
[5] https://www.sec.gov/newsroom/press-releases/2025-130-sec-issues-exemptive-order-regarding-compliance-certain-rules-under-regulation-nms
[6] https://www.linkedin.com/pulse/can-order-protection-replaced-competing-market-forces-kelvin-to-qfgte/
[7] https://www.databoiler.com/index_htm_files/DataBoiler%20SEC%20CAT%20202512.pdf

We doubt the Commission can meaningfully use CAT data in its current form to properly serve any regulatory purposes, given its **fatal flaw**[8] and **outdated design since 2012**.[9] A wide **latency** tolerance mussed up data that one cannot roll the data forward and backward to see which trade message came in first, second, third, or a hundred thousandth. In turn, it is not suitable for Best Ex, surveillance, or regulatory market monitoring purposes. CAT's current **timestamp tolerance** is set at 50+/- milliseconds and the CAT's billion-dollar price tag does not even include the cost to firms to at least upgrade Network Time Protocol [NTP] servers to a GPS time source for accuracy. Users would not notice from the CAT data how significant "initial bias" is affecting the markets.[10]

We strongly disagree with those who advocate for requiring "every counterparty to every trade must report against a unique identifier that reconciles both sides … and making CAT data public." Trade Reporting is outdated. Frequent transmittal of data in-and-out and within CAT system, unnecessary data-in-motion[11] traffic is a costly waste and more susceptible to defects. Our experience in Europe concluded that ISO-20022 is NOT suitable for low-latency, high-volume trading environment due to its verbosity and inherent latency. Embedding MMT trade lifecycle information into real-time market data feeds introduces several inefficiencies.[12] OTC bilateral activities are subjected to a deferral regime. Dual-sided reporting can never be cost justified.

The best source of trade data always resided in the clearing and settlement systems. The SEC and FINRA never lacked trade data in MIDAS and OATS in the first place. CAT was meant to augment the missing order level details. There are

---

[8] https://tabbforum.com/opinions/is-clock-synch-the-cats-fatal-flaw/

[9] https://www.linkedin.com/pulse/cat-outdated-design-since-2012-kelvin-to/

[10] Without TLE to make market data available securely in synchronized time, it causes "***initial bias***" that exacerbates the gap between subscribers of proprietary feeds and the public consolidated tape. NYSE's rooftop antenna service is a controversy, it is used by high frequency trading (HFT) firms such as Virtu, as well as McKay Brothers and its sister company Quincy Data that advocate for *"level the playing field."* NASDAQ's equalization project in removing its GPS antenna may ease conflicts among subscribers of collocation service. Yet, a directive like China banning HFTs from using client-dedicated servers co-located within exchange data center, it merely pushes the latency arms race to the nearest neighborhood. We criticize the SEC for not going far enough with the latency requirement (same manner and using the same methods) in Market Data Infrastructure Rules. Not until a mandate of market data available securely in synchronized time and the questions of who owns the data being addressed, inequitable situations would remain.

[11] https://www.databoiler.com/index_htm_files/DataBoilerInMotion.pdf ; When data is "*at-rest*" rather than "*in-use,*" it serves no value other than one must pay for storage of the data. As data is redundantly stored on original data sources, then filter through broker-dealers' systems and at the CAT system and then is regurgitated in bulk to CAT users' systems, causing significant wastages.

Instead of "*SEND*," "*OBTAIN*" or Read-Only permission to "*wiretap*" data legally (expressed consent from data owners) at its source is a substantially better approach. Wiretapping is the **fastest**, the approach would **take out the middlemen** – e.g. Trade Reporting Facilities (TRFs), and has the following advantages:

- **Benefits of Consistency** – economy of scale for centralized data management, minimize data-in-motion for cybersecurity and privacy protection, data quality is no longer a problem because what being shared is fair to everyone, avoid conflicts/ arbitrations between multiple versions of truths.
- **Prevent single point of failure** – when one TRF is down, X # of broker-dealers' data would be missing, whereas one broker-dealer's connection with SIP or CAT *intelligence layer* is down, implication is far less. When SIP is down, experience will be consistent for everyone, rather than some have the information, and some do not.
- **Values of Bespoke Model connecting to everyone** – enable the direct administration and enforcement of rights and obligations, mass customization through the powerful infrastructure, no melding nor favoritism by middlemen to distort or subjectively allocate incentives.

By no mean our suggestion inferring the Commission to recreate the bespoke model already built by vendors such as my former employer Broadridge that have API links to most if not all US broker-dealers for beneficial shareholders' information. Collecting data beyond the scope of 13(f) filings or underlying beneficial shareholder data require separate rulemaking if deviate from Congress' CAT mandates. See our response to Q2 that CAT may constitute as an unauthorized "***census***" falls outside the SEC's statutory authority.

[12] https://www.linkedin.com/pulse/addressable-liquidity-otc-equity-trading-kelvin-to-e76we

good reasons High Frequency Trading firms (HFTs) use Precision Time Protocol (PTP) instead of NTP in lining up order and trade sequence in sub-microsecond, if not in nano-second precision – this together with insights obtained from Futures and derivatives markets, plus various alternative data sources, are how they are able to see events unfolding before some of the Self-Regulatory Organizations (SROs) do.

If regulators are not seeing the same layers of intelligence as the HFTs in real-time, using the CAT to "identify the Customer responsible for market activity" is a hallucination or a delusional/ slogan-based diplomacy. CAT is retrospective (data is being reported at T+1, and accessed only at T+5), and thus it is no good to anyone other than its Cloud storage and trade reporting vendors. Its existence may merely serve an Administrative Procedure Act (APA) purpose rather than the Congress's original mandate. The Commission could have cited data from respective original sources to substantiate rulemaking, instead of using regurgitated data in the CAT centralized vault. CAT is completely blind to a **flash crash** while it is actively occurring.

Congress's original mandate to set up CAT was for the purpose of **flash crash** prevention. However, there is no mention of **flash crash** anywhere throughout the SEC's entire concept release for CAT, with the exception that SEC Director Selway did note the May 6, 2010 **flash crash**[13] in a speech at the Intermarket Surveillance Group conference.[14] Following the 2010 flash event, over 18,500 mini flash crashes have occurred in individual stocks,[15] yet the **CAT did NOT generate a single alert**, nor did it in any way avert the trend of continuously rising Ultrafast Extreme Events.

CAT should ONLY be used to capture suspicious activities from surveillance and regulatory market monitoring systems. The ONLY permissible regulatory purpose of CAT must tie-in to improvements in volatility interruption mechanisms and enforcement of the Rule 15c3-5 market access rule regarding hardcoded risk checks, credit limits, and runaway algorithm "kill switches" that live inside the brokers' and exchanges' internal servers to actively block a bad order, or halt a stock via LULD bands before a crash spreads. Any alternative use or collection of CAT data constitutes "**function creep**."[16]

2. **Are there features of the CAT that could be eliminated** because they are unnecessary or could be replaced by other existing audit trails or data sources? Identify such features and alternatives.
Anything that is NOT following proper issuance of subpoenas or NOT tied to cases being prosecuted or suspicious fraud or violation activities pertaining to "flash crash," or "market manipulations," ought to be removed from the CAT.

The SEC's articulated "scope" in the January 2026 modified version of CAT regarding Customer and Account Information System (CAIS),[17] is again a deviation from the Congress's mandate about "flash crash prevention." We respectfully disagree with the SEC's footnote 140 that said, "the CAT is designed to provide a comprehensive audit trail for U.S. securities markets, and thus is designed to provide regulators different information than what is contained within section 13(f) filings or underlying beneficial shareholder information."

An "audit trail" does NOT mean the entire universe of everyone, every order and trade activity at any time in the U.S., but ONLY relevant information pertaining to the Congress's mandate (see our response to Q1 regarding permissible regulatory purposes versus function creep). Congress conferred the authority to conduct **census** to the Department of Commerce, not the SEC. CAT in its current form is incompatible with Vice President JD Vance's remarks about

[13] https://en.wikipedia.org/wiki/2010_flash_crash
[14] https://www.sec.gov/newsroom/speeches-statements/selway-remarks-isg-conference-052726
[15] https://pmc.ncbi.nlm.nih.gov/articles/PMC5962080/
[16] Weissman v. Nat'l Ass'n of Sec. Dealers, 468 F.3d 1306, 1312 (11th Cir. 2006); Sparta Surgical Corp. v. Nat'l Ass'n of Sec. Dealers, 159 F.3d 1209, 1213 (9th Cir. 1998).
[17] https://www.sec.gov/files/rules/sro/nms/2026/34-104586.pdf

"American AI will not be co-opted into a tool for authoritarian censorship."[18] The CAT has significant privacy and security issues,[19] as well as civic concerns about **Massive Government Surveillance**.[20]

The U.S. Government Accountability Office (GAO) should investigate if CAT may be cross-subsidizing SROs in fulfillment of Exchange Act Rule 17a-1 record retention requirements and that thus deviate from the CAT's project's original purposes, given the SEC on page 49 stated "[the exemptive] relief is necessary in order to effectuate the Proposed Amendment, as Rule 17a-1 would otherwise require the customer data and information in CAIS be preserved by the [SRO] Participants."

Commissioner Peirce provided some very wise comments on CAT at the SIFMA Ops 2026, "Let's make sure that we are not collecting more than that because this has really profound implications … for Americans' privacy. When we are watching everything that you do in the markets, that is very personal to people. It is something that we typically don't like in other aspects of our lives, so I am not sure why that is okay in our financial lives. So, I think we really need to minimize what we collect. We need to think about how we protect the information that we collect, and that is something that I know that the industry has been concerned about because they feel they don't have insight into how that information is being collected..." [21]

Adding to Commissioner Peirce saying about "lawyers' billable hours", it was lobbyists' promotion of "Golden source" cloud data storage that diverted the CAT's essential focus. It ended up with, almost all resources deployed wrongly to build a giant centralized data vault, unnecessarily moving data from one place to another. CAT becomes a prime target for security and privacy hacking – i.e. honeypot risk.[22] According to a Flexera finding in 2023, "organizations lose about 28% to 32% of the money they spend on the cloud to unnecessary bloat."[23]

As much as we would like to score some quick wins in considering minor changes to the CAT that in turn would yield substantial benefits on the cost and reducing the CAT's footprint, the fundamental problem lies within unlimited desires – wanting the CAT to be "everything everywhere all at once" to everyone. Pleasing the SROs or anyone else to expand the usage of CAT data is a departure from the Congress's original mandate and it could potentially be considered an alleged misallocation of CAT funding resources.

The reluctance of the Commission and the CAT LLC Operating Committee in January 2026 approved amendment to completely purge/ eradicate non-public and personally identifiable information (PII) exacerbates systemic privacy and security risks (see Q16 part d). Then, another modified version in March 2026 projected to squeeze in more savings if the operating committee is given another $3.53 million to $5.35 million.[24] In turn, the CAT operating committee keep their jobs while the Commission staffers' can endlessly claim hours for data cleansing and conveniently point to non-transparent CAT data as "justification" for new policies, etc. There is **NO ACCOUNTABILITY** if the Commission's leadership does not restrain and realign CAT back to its original scope mandated by the Congress.

CAT serves no meaningful purpose for SROs to perform their regulatory duties in market monitoring, surveillance, and flash crash prevention (see Q1). CAT covers up the quirks that Exchanges do not want the SEC to know (see Q17) and it

[18] https://www.presidency.ucsb.edu/documents/remarks-the-vice-president-the-artificial-intelligence-action-summit-paris-france
[19] https://www.linkedin.com/pulse/privacy-security-concerns-cais-cat-consolidated-audit-kelvin-to-yqoee/
[20] https://cs.stanford.edu/people/eroberts/cs181/projects/ethics-of-surveillance/ethics.html ; https://www.eff.org/deeplinks/2023/05/10-years-after-snowden-some-things-are-better-some-were-still-fighting
[21] https://vimeo.com/showcase/12242669?video=1191677904 (from 6:36 to 10:00)
[22] https://www.linkedin.com/pulse/cat-through-z-security-privacy-requirements-kelvin-to/
[23] https://www.cloudzero.com/blog/cost-of-cloud-computing/
[24] https://www.sec.gov/files/rules/sro/nms/2026/34-105107.pdf

should no longer be structured as an NMS plan (see Q4) unless redrawing the lines of Reg. NMS and realigning rights and obligations between lit Exchanges, Wholesalers, Alternative Trading Systems (ATSs) and Single Dealer Platforms (SDPs),[25] see Q20 and Q25. The CAT is broken beyond repair and is wholly inadequate for the challenges of the twenty-first century (see Q3). Past administrations did not set it up right (see Q22). The ONLY way to reduce the CAT's footprint is right coursing its scope and architecture. We echo Commissioner Peirce's comment, [now] is the opportunity to think big thoughts about CAT, see Annex 2.

3. **Should the Commission eliminate the CAT** in favor of developing a different audit trail or data source? Why or why not? How should a new system differ? What improvements could be gained?

The time has come to defund the unsustainable CAT system. Framing its replacement as a choice between developing a different audit trail or relying on legacy data sources is a false dichotomy; far more innovative alternatives exist. Regulators and SROs do NOT need another centralized audit trails to do their jobs. Digital trails in fact already exist publicly, albeit largely in unstructured forms. The focus must shift away from collecting data as the answer; what is truly missing is the intelligent analytical layer.

Ever since the CAT was introduced over a decade ago, market participants have stressed the need to analyze the interplay between securities and CFTC-regulated Designated Contract Markets (DCMs) and Swap Execution Facilities (SEFs). The recent memorandum of understanding between the SEC and the National Futures Association (NFA) is a step in the right direction.[26] The SEC-CFTC harmonization initiative maintains existing statutory boundaries while collaborating to unify cross-market analysis.[27]

Multi-asset convergence between TradFi and DeFi is inevitable and new risks have emerged. For example, the SEC approved $DOJE is a crypto Exchange Traded Fund (ETF) with underlying "assets" being MEME Coin, despite a statement by the SEC Division of Corporate Finance on Feb 27, 2025[28] which stated that "MEME Coins for entertainment and social cultural purposes are NOT securities." Also, any "Digital Collectables" that do not involve CEA regulated commodity options, futures, or leveraged OTC transactions would be outside scope of CFTC's oversight.[29]

Reliance on the regulated structure of the product and the transparent disclosures for the valuations and the regulatory controls may be an insurmountable reality. That guardrails may merely displace fraud and manipulation risk from SEC-regulated markets, rather than preventing that risk from proliferating elsewhere. Bad actors / foreign adversaries play across markets and payment systems simultaneously. All five Dodd-Frank regulatory agencies (SEC, CFTC, OCC, FRB, FDIC) must break down silos and work in unison to safeguard US financial stability.

Regulatory market monitoring ought to be modernized alongside market and technology evolutions (Agentic Overlay). The goal is to assess changes in market dynamics, understand where frictions and liquidity concentration may shift at rapid pace or irregularities and dysfunctional market behaviors occur. Analyze the stress points and weak links where

[25] https://www.databoiler.com/index_htm_files/DataBoiler%20BIG%20OPP.pdf
[26] https://www.sec.gov/files/sec-nfa-2026.pdf
[27] https://www.isda.org/2025/10/02/a-path-to-greater-cftc-sec-alignment/ ; https://www.sec.gov/files/isda-sifma-letter-cftc-sec-harmonization-letter-051926.pdf ; https://www.sec.gov/newsroom/speeches-statements/selway-remarks-global-exchange-fintech-conference-060426
[28] https://www.sec.gov/newsroom/speeches-statements/staff-statement-meme-coins
[29] https://www.databoiler.com/index_htm_files/DataBoiler%20CFTC%2020250818.pdf

the next **flash crash** or systemic risks may emerge, and develop mitigation protocols with more effective and timely volatility interruption mechanisms.[30]

While the Commission and SRO staffers understandably worry about AI displacement, they must learn to orchestrate these new tools. *A good decision, made now and pursued aggressively, is superior to a perfect decision made too late.* Humans are slow and CANNOT manually reconcile the massive volume of structured trade logs and unstructured data driving modern markets. AI bridges this gap by handling tedious data ingestion and synthesis. Supported by human context, institutional knowledge, and strict guardrails, AI acts as a force multiplier – not job replacement. This shift elevates agency personnel from manual data processors to strategic gatekeepers of market integrity.

See Annex 2 in this comment letter for an elaborated discussion of our innovative design to make CAT replacement nimble and effective with **agentic AI**, Model Context Protocol (**MCP**) to connect and orchestrate workflows across systems, Retrieval-Augmented Generation (**RAG**) for knowledge retrieval to avoid wasting tokens on hallucinations, and surround both with **Zero-Trust** cyber defense (see Q28) to safely manage what the AI can see and do. Collectively these components form the "layers of intelligence" to address challenges of today and the future.

[30] NOTE: **Volatility interruption mechanisms** (e.g., replacing Single-Stock Circuit Breakers with the LULD Mechanism; overhauling Market-Wide Circuit Breakers; banning Stub Quotes and regulating Market Maker obligations; standardizing *"Clearly Erroneous"* trade executions) were piloted or implemented before the Congress' CAT mandate in 2012. The subsequent transitioning from *"pilots"* to permanent laws was a natural progression. Credits should be given to Reg. SCI that mandates all major proprietary trading firms and broker-dealers to participate in at least one coordinated cross-market Level 1 and Level 2 circuit breaker test per year.

Section B — Structure and Governance of the CAT

4. **What are the advantages and disadvantages of structuring the CAT as an NMS plan?** Should it continue to be structured this way?

CAT should no longer be structured as an NMS plan, unless redrawing the lines of Reg. NMS and realigning rights and obligations between lit Exchanges, Wholesalers, ATSs, and SDPs (see Q20 and Q25).

We are suspicious about the focus of unique identifiers – CAT Customer-ID (CCID), and how it may be misused rather than fulfilling the Congress's original mandate for flash crash prevention (see our response to Q1). CAT data may be skewed in undermining non-lit trading venues' performance and in covering up known issues on lit Exchanges. Empirical researches have proven that "exchanges optimally restricted access to price information",[31] and "how Trade Reporting Facilities (TRFs) handle and preference orders can impact execution quality."[32] Using an NMS plan to build CAT creates deep **conflicts of interest** and shields individual actors:

- **Diffusion of Responsibility:** When everyone is collectively responsible, no single entity takes ownership. SROs can hide behind the "Operating Committee" voting block to stall changes, mask operational failures, or pass massive cost overruns down to broker-dealers and retail investors.
- **The "For-Profit" Conflict:** Modern exchanges are publicly traded corporations trying to maximize trade volumes. Giving a committee of for-profit competitors the joint duty to manage a public surveillance utility turns the CAT into a politicized battleground over funding models, rather than a neutral, agile data tool.
- **Regulatory Cost Bloat:** Because it is structured as an insulated NMS plan, the CAT became an unchecked "money pit." It took sweeping interventions to slash data-retention scopes and cap runaway budgets because the SROs could not efficiently govern themselves.

It is absolutely opaque regarding why FINRA CAT LLC relied on a **closed-door contract** of nearly **$770,000 annually** to pay for a so-called "Market data vendor fees."[33]

What data elements mandated under §6.5(a)(ii) of the CAT NMS Plan – including SIP data, OPRA options data, and LULD price bands—are not already possessed by the SROs, such that Algoseek must provide them?

What justification does it have to fund a slow, software-based formatting loop hosted entirely on standard AWS cloud servers, instead of leveraging established market data vendors whose precision infrastructure's (e.g. FPGA, native PTP time synchronization, calibration / normalization / ingestion of ultra-low latency feeds) costs are shared globally?

What actionable data elements or specialized processing does this "market data vendor" offer that could not be achieved more cost-effectively and with greater scalability through native AWS infrastructure?

How is it NOT a structurally redundant data loop if it is a basic CPU-bound data-wrapping process to regurgitate SIP and other public data mentioned earlier? How is it NOT introducing noise or distorting the NBBO sequencing of orders and trades in CAT when it serves no useful calibration or technical normalization purpose beyond fitting stale data?

SROs already generate comprehensive Level 3 depth-of-book proprietary feeds. By failing to modernize market data infrastructure of the public consolidated tape, for-profit SROs exacerbated structural latency and content gaps to

[31] https://www.bayes.citystgeorges.ac.uk/__data/assets/pdf_file/0011/366599/sale-price-information-cass-knowledge.pdf
[32] https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3369025
[33] https://www.sec.gov/files/rules/sro/ise/2025/34-102210-19b-4.pdf ; https://www.catnmsplan.com/sites/default/files/2026-04/03.31.26-CAT-2026-Budget.pdf

protect their lucrative proprietary data and collocation monopolies.[34] CAT LLC could have directly aggregated these high-speed, native pipelines to augment the missing order-level details that legacy systems like OATS and MIDAS lacked. Had regulators simply synthesized these existing proprietary feeds alongside centralized clearing and settlement systems, the SEC and SROs could have conducted comprehensive market-wide analysis without ever needing to construct a centralized, multi-billion-dollar CAT data vault.

NMS Plan is meant to facilitate harmonization among SROs, for example, to avoid Limit Up-Limit Down (LULD) price band issues. The August 24, 2015, market sell-off reflected non-harmonized logics at NYSE and NASDAQ that caused back-to-back halts.[35] This observable "double halts" was not hidden data requiring CAT to uncover. Also, the upgrade of NASDAQ's LULD cross price protections[36] are regulatory responses to known phenomena – not prompted by CAT.

Sadly, the CAT NMS plan forces the entire financial industry to redundantly report data at an astronomical cost, exploiting mandatory regulatory fees to fund an unamortized cloud pipeline. The Commission disapproved CAT limited liability provisions for SROs in 2021.[37] No enforcement action has ever been brought against CAT LLC, despite the fact that Reg. NMS intends to treat all national stock exchanges, options exchanges, and FINRA as a single entity for collective liability.

Widespread industry and public frustration over CAT governance threatens to destroy faith in Reg. NMS and the broader U.S. securities markets. Reg. NMS advantages are NOT realized especially when the SEC continues to tolerate chronic national market disorders. The SEC, GAO, and the U.S. Department of Justice (DOJ) Antitrust Division must immediately investigate uncompetitive structural deficiencies, as well as assess whether CAT data was **weaponized for political or commercial reasons** that benefit the SROs and their vendors.

5. **If not an NMS plan, how should the Commission and/or SROs direct and oversee the CAT?** What benefits, actions, and costs would be associated with transitioning? Would benefits offset costs?

   The Commission and the SROs have repeatedly let down the public interest for failing to directly and indirectly oversee the CAT. CAT's original funding and subsequent executed share models are inequitable, with multiple lawsuits filed

[34] Unfairness occurred because of initial bias (see footnote 10). Average investors would not be aware if they receive inferior price or not, given SROs control the SIP NBBO and most retail order flow is interacted with HFTs instead of Exchanges. In essence, HFTs purchase most of Exchanges' production capacity and prioritize processing given their latency advantage, then HFTs turnaround to sell any excess bandwidth they do not use to generate alpha in providing outsource execution services to cover their ultra-low latency infrastructure costs, amid paying a portion of their super-tier rebates from Exchanges as payment for order flow to subsidize zero-commission offered by retail brokers. The implementation of Market Data Infrastructure Rules (MDIR) has not yet seen results in lowering market data cost and "*ensuring that* [SIP's] *core data evolves along with the broader market ecosystem*."

[35] https://www.sec.gov/newsroom/speeches-statements/2013-ts051613mjwhtm The issue of varying logic causing back-to-back halts at different SROs was not hidden data that required CAT to uncover. August 24, 2015 market sell-off is a public phenomenon where ETFs and individual stocks suffered over 1300 LULD halts Because exchanges used mismatched rules to resume trading, many stocks reopened, instantly printed a highly volatile price against a lingering order imbalance, and triggered a second 5-minute halt just milliseconds later. Everyone could see this event unfold in real-time on standard market data feeds without the need of CAT.

[36] https://listingcenter.nasdaq.com/rulebook/NASDAQ/rulefilings The industry already knew that if a stock was locked in a limit state right before the close, the extreme concentration of Market-on-Close and Limit-on-Close orders could trigger artificial price spikes or crashes when forced into a single closing cross execution. The upgrade was based on the known mechanics of NASDAQ matching engine to shield the closing cross from late-day algorithmic cascades, not because a CAT query alerted them to a new behavior.

[37] https://www.sec.gov/files/rules/sro/nms/2021/34-93484.pdf

against it.[38] It is no longer about cost-benefits, but salvaging the reputation of the Commission, FINRA, and other SROs. We pointed out these 4 key points in the past,[39] which are still valid today:

(a) If the CAT fee is related to supporting the SEC to "rapidly reconstruct market events/ trading activity" beyond using the public available data, then the Commission could alternatively subscribe to the SROs' proprietary feeds for any non-public data, or seek expressed consent to voluntarily share, or use of its permissible authority to summon the relevant private information.

(b) If the CAT fee is related to "facilitating risk-based examinations" and/or "improving abilities for evaluating tips, complaints and referrals of potential misconduct made to regulators, monitoring and evaluating changes to market structure," then the SEC and SROs should go back to the Congress for funding or pay for it using collected fines, penalties, and intragovernmental fees, but not rely on "user fees" which hurt broker-dealers and retail investors who have to pay it.

(c) If the CAT fee is related to "better identification of potentially manipulative trading activity, increased efficiency of cross-market and principal order surveillance," then private surveillance businesses affiliated with Exchange Groups that stand to receive benefits from the CAT should therefore pay the most if not all of such CAT costs. The SEC and other SROs shall have the choice to use peers' surveillance system or build their own or buy from other private vendors rather than relying on the CAT.

(d) If the CAT fee is related to "improving efficiencies from a potential reduction in disparate reporting requirements and data requests," then it should be segregated into the regulators' portion and the users' portion. If CAT is constituted as one of the "user fees" imposed by the SEC and/or SROs, then according to the GAO, these "fees assessed to users for goods or services provided by the Federal Government are deposited to the Treasury as miscellaneous receipts and are generally not available to the agency." Fees assessed under the authority of the Independent Offices Appropriation Act of 1952 (codified at **31 U.S.C. § 9701**), rather than under a specific authorizing statute, must be deposited to the Treasury as miscellaneous receipts and are not available to the agency or program that collected the fees, unless otherwise authorized by law.

To restore public faith (see our response to Q4), it is time to defund and dismantle CAT, decommission the CAT operating committee, conduct independent investigations, revisit the CAT Design and Scope with a fresh perspective (see our response to Q3 and Annex 2).

6. **Should the Commission amend the CAT NMS Plan to implement a different voting structure?** What should it be and why? Should affiliated SRO groups or non-affiliated SROs receive extra votes?

When the D.C. Circuit Court of Appeals struck down the SEC's MDIR Governance Order—which had attempted to dilute exchange power by establishing a 1/3 non-SRO and 2/3 SRO voting structure—there is statutory limit under the Exchange Act that the SEC cannot force SROs to share their regulatory voting power with private commercial entities

[38] https://www.databoiler.com/index_htm_files/DataBoiler%20SEC%20CAT%20Funding%20202108.pdf ; https://www.databoiler.com/index_htm_files/DataBoiler%20SEC%20CAT%20Funding%20202212.pdf ; https://www.sidley.com/en/insights/newsupdates/2025/08/eleventh-circuit-vacates-secs-2023-funding-order-for-the-consolidated-audit-trail-cat ; https://www.jonesday.com/en/practices/experience/2025/07/citadel-securities-wins-eleventh-circuit-decision-invalidating-sec-order-creating-funding-mechanism-for-cat ; https://www.sifma.org/advocacy/letters/prohibition-on-the-use-of-reserve-funds-by-the-consolidated-audit-trail-national-market-system-plan ; https://www.americansecurities.org/post/asa-citadel-securities-file-lawsuit-challenging-sec-s-consolidated-audit-trail-funding-order ;

[39] https://www.linkedin.com/pulse/cat-bifurcated-cost-allocation-inequitable-kelvin-to/ ; https://www.linkedin.com/pulse/cat-both-original-funding-executed-share-models-inequitable-kelvin-to/

that do not bear those same statutory liabilities.[40] Why would anyone believe the current administration of the SEC can amend the CAT NMS Plan to implement a different voting structure? Attempt to allocate one vote per affiliated exchange group (rather than per license) is insufficient to avert conflicts of interest and other disadvantages as discussed in our response to Q4. It is too late to attempt to amend the CAT NMS Plan. The CAT is broken beyond repair and is wholly inadequate for the challenges of the twenty-first century (see our response to Q3).

7. **Should the Commission amend the CAT NMS Plan to require a different vote threshold** (majority, supermajority, unanimous) for any specific actions? Which actions and why?
   Not useful, see our response to Q6.

8. **Are there measures that could increase transparency and accountability** around CAT NMS Plan voting while protecting sensitive information? Should more information be made public?
   Nothing around CAT NMS Plan voting amendment would restore public trust, see our response to Q6. The ONLY way to improve transparency and accountability is by the SEC, GAO, and DOJ Antitrust Division's immediate investigations to determine if CAT has any uncompetitive structural deficiencies, as well as assess whether CAT data was **weaponized for political or commercial reasons** that benefit the SROs and their vendors. Please see our response to Q4.

Advisory Committee

9. **What powers, responsibilities, or rights should be given to the CAT Advisory Committee?** How would these affect SRO governance? Would they raise conflicts of interest?
   We do NOT believe any CAT Advisory Committee can do anything meaningful to avert SROs' **conflicts of interest** and other disadvantages under NMS Plan as discussed in our response to Q4.

   Few in the industry dare to say anything against SROs because of worry about retaliation. The SEC maintained a formal representative observer seat on the CAT Advisory Committee and has been embedded in the operational process since Rule 613 was first finalized. Despite the authorities' close oversight, the project migrated away from its explicit, post-2010 Flash Crash mandate – which was to create an anonymized, cross-market forensic tool – and instead evolved into a monolithic, un-vetted data repository. It is vulnerable to security threats, intruding upon privacy and impairing the civil liberties of Americans who are "transacting" or directly or indirectly "engaging" in any way, shape, or form in the U.S. securities markets, see Q26.

   The SEC's MDIR Governance Order attempted to provide additional power, responsibilities, or rights to non-SROs Advisory Committee (broker-dealers or asset managers). It demonstrated how an Advisory Committee is a toothless Public Relation shield and is ineffective in bringing forth changes (e.g. cannot veto a budget, they cannot alter data collection, and they cannot block a security protocol). D.C. Circuit Court held that the SEC cannot force SROs to share their regulatory voting power with private commercial entities that do not bear those same statutory liabilities (see our response to Q6).

   When the SROs terminated the Plan Processor contract with Thesys in 2019 and selected FINRA CAT LLC, the industry was promised a "fresh start". Replacing the Plan Processor while keeping the exact same structural blueprint is like changing the pilot of a plane that was built without wings. It guarantees a crash, regardless of who is in the cockpit. CAT is fundamentally flawed by architecture, broken beyond repair and is wholly inadequate for the challenges of the twenty-first century (see our response to Q3).

[40] https://law.justia.com/cases/federal/appellate-courts/cadc/21-1167/21-1167-2022-07-05.html

10. **Should membership of the Advisory Committee be reserved for certain interests or individuals with specific experience?** Should it be expanded to include ATS operators, technology experts, or others? How should representation be apportioned?

We at Data Boiler are uniquely qualified to give advice on CAT. We are designated by the European Commission as member of the Data Expert Group. See Annex 2 to review our recommended architectural redesign. Please keep us posted where our expertise might be helpful for the CAT advisory committee.

Meanwhile, the current CAT advisory committee comprises high caliber members, such as Dr. Gregg Berman (former Associate Director at the SEC that led the Thesys built MIDAS project and is currently a Director of Market Analytics and Regulatory Structure at Citadel Securities). Reference to Q22, the Commission's intention to modify the "representative order linkage requirements" is a tacit admission that the original "daisy chain" approach supported by Dr. Berman was a flawed, incredibly expensive dead end.

The question is NOT about who is who on the advisory committee, or who was, or is, the CAT Plan Processor that will or will not listen to advice, but rather the faulty and outdated design of the CAT itself. Again, we would like to reiterate – CAT LLC could have directly aggregated high-speed, native pipelines to augment the missing order-level details that legacy systems like OATS and MIDAS lacked. Had regulators simply synthesized these existing proprietary feeds alongside centralized clearing and settlement systems, the SEC and SROs could have conducted comprehensive market-wide analysis without ever needing to construct a centralized, multi-billion-dollar CAT data vault (see our response to Q4).

That for-profit SROs fiercely defended their lucrative stakes is understandable. However, attending meetings and providing views does not mean the SROs CAT operating committee would listen, cooperate, and not take matters to Court. We have reservations that the CAT Advisory Committee can make meaningful changes to right the course of the CAT's scope and design in addressing its fundamental flaws. Please also see our response to Q9.

11. **Should the Advisory Committee have additional rights** beyond attending meetings and providing views? If so, what rights and why?

Even though the SEC maintained a formal representative observer seat on the CAT Advisory Committee and has been embedded in the operational process since Rule 613 was first finalized, we doubt the SROs would cooperate. More likely than not, disagreements end-up in Court costing taxpayers to foot the bills.

Section C — CAT Funding and Cost Management

12. **Should the Commission amend the CAT NMS Plan to require additional cost-management measures?** What measures and why?

This CAT money pit is already unsustainable and will only be exacerbated. The continuous rise in trading activities that derives from U.S. advantages over other jurisdictions, openness to new ideas, tenacity to accelerate market developments, including but not limited to creation of new trading venues, extended trading hours, and expanded range of securities products (e.g. 0DTE options, crypto ETPs). Consider these opportunities approved and endorsed by a mix of Congressional action (e.g. ancillary assets)[41] and the Commission approved policies. Nothing can defy the growing volume of data from various flood stream sources flowing in the same direction like a "tsunami" and crushing a "reservoir" or "centralized vault".

Had the CAT operating committee tried to put cost control measures? Yes, they do provide their budget, albeit for example it is absolutely opaque regarding why FINRA CAT LLC relied on a **closed-door contract** of nearly **$770,000 annually** to pay for a so-called "Market data vendor fees," see our response to Q4. Cloud hosting services is 60+% of the $128.6 million annual total technology costs, legal cost is an eyepopping near $9 million bill (no wonder Commissioner Peirce teased the CAT project by stating "We have set it up in a way that has really been very good for lawyers. There have been, for years now, people who have been getting lots of billable hours"),[21] insurance takes $1.25 million ($100M insurance cap grossly undermines a National security threats – a breach is not a minor corporate loss; it could trigger a structural collapse of U.S. capital markets), etc.[33] We do NOT believe control measures, such as monitoring and metering the ad-hoc queries of enormous data sets would move the needle in shaving "fat" off this "outsized elephant."

Referring to our comment letter in May 2021,[42] "an eye-popping $135 million a year for cloud hosting services… this number is 20 to 100 times more expensive than private sector metrics ..." A Trader Magazine article by Kirsten Wegner at Modern Markets Initiative reiterated the cost concerns in 2024,[43] "CAT … Most Americans Have Never Heard … has a budget that dwarfs that of many federal agencies, and its ever-ballooning costs threaten to eat up the returns of your investments and savings… these costs have recently been approved to be passed through to and paid by U.S. stock market participants…" Given that CAT's annual operating budget hovers between $188 million and $272 million,[44] its funding now exceeds the individual annual budgets of several federal agencies.

The Jan 2026 CAT NMS Plan amendment squeezed approximate $7-9 million in annual cost savings which is a tiny fraction and nowhere in proportion with the CAT's budget. Then, the March 2026 modified version suggests deeper cuts aiming to generate $50 to $70 million in annual savings, but only if they incur an additional $3.53 to $5.35 million in one-time implementation of the following changes:

- Interim CAT Order-ID Amendment (one-time implementation cost $200,000 – $400,000);
- Data Storage Amendment (one-time implementation cost $165,000 – $265,000);
- Late Data Re-Processing Amendment (one-time implementation cost $250,000 – $500,000);
- Options Market Maker Quotes Amendment (one-time implementation cost $135,000);
- Rejected Message Amendment (one-time implementation cost $75,000 – $150,000);

[41] https://www.databoiler.com/index_htm_files/DataBoiler%20CFTC%2020250818.pdf
[42] https://www.databoiler.com/index_htm_files/DataBoiler%20SEC%20CAT%2020210503.pdf
[43] https://www.tradersmagazine.com/am/cat-how-can-a-new-government-program-that-most-americans-have-never-heard-of-cost-more-than-the-annual-budget-of-half-the-federal-agencies/
[44] https://www.catnmsplan.com/sites/default/files/2024-11/11.20.24-CAT-LLC-2025-Financial_and_Operating-Budget.pdf

- Data Availability Amendment (one-time implementation cost $200,000 – $400,000):
- Reference Data Amendment (one-time implementation cost $2.5 – $3.5 million)

None of these would be necessary if FINRA CAT LLC could have directly aggregated high-speed, native pipelines to augment the missing order-level details that legacy systems like OATS and MIDAS lacked. Had regulators simply synthesized these existing proprietary feeds alongside centralized clearing and settlement systems, the SEC and SROs could have conducted comprehensive market-wide analysis without ever needing to construct a centralized, multi-billion-dollar CAT data vault. Sadly, the CAT NMS plan forces the entire financial industry to redundantly report data at an astronomical cost, exploiting mandatory regulatory fees to fund an unamortized cloud pipeline.

The billions of dollars poured into the CAT project — alongside extensive industrywide efforts in quote and trade reporting — represents a near-complete waste. The system has failed to achieve its regulatory objective (see our response to Q1), despite FINRA CAT LLC taking over operations from Thesys in 2019. Widespread industry and public frustration over CAT governance threatens to destroy faith in Reg. NMS and the broader U.S. securities markets.

If the Commission decides to absorb CAT cost rather than unjustly passing through these CAT costs to US stock market participants, taxpayers would be uneasy with a further ballooning of the SEC's $2.15 billion budget by another 7 to 8%. Self-serving or serving the public? Also, see our response to Q23. Choose wisely and do not procrastinate to stop this money pit. Modernize the execution of its original regulatory mandate to tackle new challenges (see Q3 and Annex 2).

13. **Should the Commission require the SROs to adopt a different funding model** for the CAT? What model and why?

Regarding the CAT's funding, both the original tiered fee structure and subsequent revised Executed Share Model are like a **Financial Transaction Tax**.[45] The Eleventh Circuit Court of Appeals vacating the 2023 Funding Order. Rather than tying regulatory fees to the actual recipients of its public benefits, the CAT operating committee (composed exclusively of SRO representatives, who hold concentrated, unchecked power to establish budgets and collect fees) has been able to shift the cost burden to the non-SRO entities. The financial arrangement when FINRA CAT LLC took over the CAT Plan Processor role from Thesys in 2019 has caused *immense inequalities* in funding by SROs attempting to *"recoup"* hundreds of millions in development costs that are actually **private assets**, while forcing regular industry members to *foot the bill* for AWS cloud services and lucrative vendor legal fees.

Our past analysis demonstrated that the CAT cost-allocation framework creates an inequitable ecosystem, disproportionately *penalizing smaller firms* while shielding the industry's largest participants.[46] The top 36 market players generate an astounding 96.67% of all message traffic, in the process receiving deep discounts, maximum caps, and special rebate treatment. Meanwhile, a massive 97.1% of smaller broker-dealers generate a mere 3.33% of message traffic but are hit with disproportionately high quarterly minimums. This creates a *massive administrative wastage* in billing and collections for *"de minimis"* fees, converting the funding model into an anti-competitive barrier to entry that suppresses capital formation.

The sheer technical design of the CAT has morphed into a *bottomless data warehouse that ballooned* from an historic average of 296 billion to a recent average of 700 to 800+ billion records every single day. This massive *data-in-motion* accumulation **cannot be economically justified**, resulting in an incredibly low capitalization rate where only a tiny fraction of expenditures is counted as actual assets. As a result, the entire industry is stuck financing a stagnant state-run monopoly data vault that *would never be approved in the private sector*. Because the authorities refuse to, in a

[45] https://securitytraders.org/wp-content/uploads/STA-FTT-Letter-FINAL-03_16_2021.pdf
[46] https://www.linkedin.com/pulse/cat-bifurcated-cost-allocation-inequitable-kelvin-to/ ; https://www.linkedin.com/pulse/cat-both-original-funding-executed-share-models-inequitable-kelvin-to/

timely fashion, rectify these baseline architectural flaws, the CAT continues to serve as an unconstrained cloud-hosting *money pit*.

Rulemaking intending to seek sole benefits for government agencies or affiliated SROs must be prohibited, yet the CAT committee continually uses high-level regulatory platitudes to expand its administrative footprint, see Q1, Q16, Q20, and Q24. The SEC and SROs *move the goalposts* under the vague guise of *"enhanced oversight"* to bypass the U.S. Constitution's separation of powers. If SROs insist on absolute discretionary control over fee setting and dispute resolutions, they must strip themselves of arbitral/ prosecutorial *immunity* so market participants and the public can utilize *antitrust laws* to fight back.

CAT was given 10+ years as an experiment, not once (Thesys), but twice (FINRA CAT LLC), to do a $2+ billion proof-of-concept about a "centralized single source of truth" that is doomed to failure. Write it off. Most if not all CAT costs are sunk cost. Little of the technology development work was capitalized. Revert to OATS or simply keep the OAT portion of post-trade data if the SEC likes to retain the "CAT" shameful name even if it has little to no substance.

Again, Trade Reporting is outdated. Frequent transmittal of data in-and-out and within the CAT system, unnecessary data-in-motion traffic is a waste of resources and more susceptible to defects.[11] The best source of trade data always has resided in clearing and settlement systems (see Q16). The industry never had complaints about OATS, COATS, and EBS in contrast to CAT because it was not onerous. Now that CAT's funding model has been struck down by the Court, the model is unsustainable. It is better to sunset the centralized CAT, and redeploy CAT resources into assembling layers of intelligence (see our response to Q3 and Annex 2).

14. **Should the Commission require the SROs to maintain reserve funds** for CAT operations? What size and purpose should such reserves serve?

The purpose of maintaining a reserve fund under the CAT NMS Plan is to ensure the financial liquidity, operational continuity, and baseline stability of this enormous database. According to the CAT NMS Plan, the CAT Operating Committee is mandated to establish and maintain an operational reserve capped at not more than 25% of the annual budgeted CAT expenses. This capital buffer is explicitly designated to serve several key operational and defensive administrative functions: insulating against cash-flow volatility; a financial Lifeline during legal warfare; funding patch work and security mandates; and preventing "Over-Collection" profit laundering.

While the original stated purpose is honorable, the reserve has essentially become a subsidized cushion that protects the SROs from the immediate financial consequences of their own architectural overruns in the current reality. With the reserve on track to be completely exhausted by August 2026, the industry is actively fighting the SROs' emergency proposals to "top it up."[47] SIFMA recommends holding reserves in escrow until a new funding path is established.

Where does the money come from for this reserve funds, do SROs grow money on trees? Is the $23.5 million in estimated liquidity reserve balance for 2026 insufficient for the orderly liquidation of this "outsized elephant"?! There is no point in topping up the reserve when the CAT is burning hundreds of millions year-over-year and is fiscally unsustainable. Humans stop hallucinating. Consider the use of Agentic AI to replace CAT.

[47] https://www.sifma.org/wp-content/uploads/2025/12/CAT-Reserve-Funds.pdf ; The 11th Circuit's decision eliminated CAT LLC's funding authority under the 2023 CAT funding model. The Court's decision prohibits CAT LLC from using any funds, including any reserves, collected pursuant to the vacated funding model to fund ongoing and future CAT costs.

15. **Should Section 31 fees or other alternative funding methods be used** to fund the CAT? What are the advantages and disadvantages?

Per our prior submitted comment letters in 2021 and 2022,[48] we strongly oppose the use of Section 31 fees. Merely pointing to the high-level purpose of CAT – "creating enhanced oversight of the markets" and attempts to claim it as "reason" is NOT acceptable. Moving the goalposts wherever they want under the guise of "enhanced oversight" undermines the Constitution that divided the Government into three branches: Legislative, Executive, and Judicial. CAT LLC's argument citing "Industry Members have far greater financial resources than the Participants" is NOT relevant. The "sh*t hit the fan" **Financial Transaction Tax** which dun everyone in the industry, and then is passed-down to the end-investors through price adjustments, is thus **unjust**.

According to SIFMA, total private-sector compliance expenditures directly incurred by industry members to report data has exceeded $1.7 billion annually.[49] Industry members have already paid more than their fair share for the CAT. Cost-benefits analysis should NOT ignore this significant social cost incurred by industry members. Section 31 like fees must be discouraged and avoided. Rulemaking to seek sole benefit for a government agency or the affiliated SROs should be prohibited because it conflicts with serving the public interest.

OATS, which ran from 1998 to 2021 never had such an unsustainable funding problem and it was never funded by a pass-through regulatory fee. Instead, OATS was built, operated, and maintained natively by FINRA as part of its general operating budget. The industry never filed lawsuits against the OATS funding structure. As another example, MIDAS is entirely federally funded through the SEC's formal technology and procurement budget, which is approved annually by Congress – a stark contrast to the CAT's closed-door contracts. It is all about mandate matching the funding source, see our response to Q5.

P.S. EBS was formally established in 1989 under the SEC's authority and later codified under SEC Rule 17a-25 in 2001 and is still running seamlessly today. EBS operates strictly as an on-demand, internal operational expense, see Q24. COATS implemented beginning in 2000 by a coalition of the major U.S. options exchange. COATS was financed directly out of the general corporate operating budgets of the options exchanges. The exchanges absorbed the localized data-ingestion costs as part of their standard regulatory overhead, funded by their native business revenues (such as listing fees, transaction matching fees, and proprietary market data sales). The structural peace surrounding EBS and COATS – No Centralized Corporate Welfare; No Usurpation of Fee-Setting Power; Proportionality and the Absence of Red Tape – highlights why the CAT funding model triggered massive litigation.

If CAT is meant to prevent future *flash crashes*, curb suspicious trading behavior and unusual market events, then why should those who are doing things fairly and squarely be subjected to regulatory scrutiny and the CAT cost burden?[50] We think those who *"operate at the edge"* and have higher risks for potential conflicts of interest – including the SROs (see Q4, Q6, and Q9), should bear much of the CAT cost given the extra efforts in deciphering their complex activities as compared to firms with a simpler business model.

[48] https://www.databoiler.com/index_htm_files/DataBoiler%20SEC%20CAT%20Funding%20202212.pdf
[49] https://www.sifma.org/wp-content/uploads/2025/06/SIFMA-CAT-Reform-Recommendations-6-6-2025-Final.pdf
[50] https://www.linkedin.com/pulse/cat-how-much-should-finra-trading-venues-pay-kelvin-to/

Section D — CAT Design and Scope

16. **Should the Commission amend the CAT NMS Plan to modify the scope of CAT data collection?** What modifications and why?

YES, **the scope of CAT must change** whether through amendment of the CAT NMS Plan or other means (e.g. updated legislation). CAT serves no relevant regulatory purpose for the SEC and the SROs to do their jobs. Despite the authorities' close oversight, the project migrated away from its explicit, post-2010 **Flash Crash** mandate – which was to create an anonymized, cross-market forensic tool – and instead evolved into a monolithic, un-vetted data repository. It is vulnerable to security threats, intruding upon privacy and impairing the civil liberties of Americans who are "transacting" or "engaging" in any way, shape, or form in the U.S. securities markets. Please see Q1 and Q2).

An "audit trail" does NOT mean the entire universe of everyone, every order and trade activity at any time in the U.S., but ONLY relevant information pertaining to the Congress's mandate. Congress conferred the authority to conduct **census** to the Department of Commerce, not the SEC. CAT in its current form is incompatible with Vice President JD Vance's remarks about "American AI will not be co-opted into a tool for authoritarian censorship." CAT has significant privacy and security issues, as well as civic concerns about **Massive Government Surveillance**.[20]

CAT should ONLY be used to capture suspicious activities from surveillance and regulatory market monitoring systems. The ONLY permissible regulatory purpose of CAT must tie to improvements of volatility interruption mechanisms and enforcement of Rule 15c3-5 market access rule regarding hardcoded risk checks, credit limits, and runaway algorithm "kill switches" that live inside the brokers' and exchanges' internal servers or actively block a bad order or halt a stock via LULD bands before a crash spreads. Any alternative use or collection of CAT data constitutes "**function creep**."

Replacing the Plan Processor while keeping the exact same structural blueprint is like changing the pilot of a plane that was built without wings. It guarantees a crash, regardless of who is in the cockpit. CAT is fundamentally flawed by architecture, broken beyond repair, and is wholly inadequate for the challenges of the twenty-first century financial markets. Digital trails already exist publicly, albeit largely in unstructured forms. **The focus must shift away from collecting data** as the answer; what is truly missing is the intelligent analytical layer. Regulatory market monitoring ought to be modernized alongside markets and technologies evolutions (e.g. multi-asset convergence between TradFi and DeFi + Agentic Overlay) to meet challenges of today and the future (see our response to Q3).

Bad actors / foreign adversaries play across markets and payment systems simultaneously. All five Dodd-Frank regulatory agencies (SEC, CFTC, OCC, FRB, FDIC) must break down silos and work in unison to safeguard US financial stability. To achieve the goals in assessing changes in market dynamics, understand where frictions and liquidity concentration may shift at rapid pace. Consider stress and weak links where the next **flash crash** or systemic risks may emerge in order to develop mitigation protocols with more effective and timely volatility interruption mechanisms. We suggest redefining scope of CAT (or a totally new **forensic tool**) as follows:

a. Covered Securities (Product Scope)

Currently Rule 613 applies to secondary market transactions in all NMS securities (i.e. all listed equity securities traded across U.S. exchanges; all standard, exchange-traded equity and index options; and OTC equity securities).

Ever since the CAT was introduced over a decade ago, market participants have stressed the need for futures, swaps, and securities-based swaps. For example, if an institutional investor or algorithmic desk creates a synthetic future by simultaneously buying a call option and selling a put option on the same underlying stock or stock index (with the exact same strike price and expiration date), this strategy is 100% within the current scope of the CAT. If a firm creates a synthetic future using over-the-counter (OTC) derivatives—such as structured single-name total return equity swaps

or forward contracts—the strategy is completely outside the scope of the CAT. Financial alchemy in itself is neutral depending on how a product is used. Agentic AI acts like the scanners in modern cars, it helps eliminate blind spots.

NOTE: Non-standard exchange-listed options are currently excluded from CAT. We understand that Corporate Action or "Adjusted" Options (The "NS" Flag) is unlikely to trigger a flash Ultrafast Extreme Event. However, there is explosive growth in the popularity of "FLEX Options", given the mass customization trend.[51] Unlike securities markets defined by continuous price discovery, trading in Prediction Markets function as short-lived exposures to the "cause", while 0DTE Options provide exposure to the "effect".[52] The beauty of simultaneous hedging strategies is the decoupling of Event and Price Risk. These critical phenomena are reshaping the financial landscape, and policymakers cannot afford to overlook their overarching convergence risks.

Exuberance rises as financial markets get caught in the crossfire of currency and trade wars.[53] A rebellious move by an insurgent with a war chest to orchestrate a market wide shake-up, and foreign adversaries wanting to erode the US's prominent market position.[54] In turn, correlations have broken between equities and fixed income markets. Traditional corporate bonds, municipal debt, and asset-backed securities are currently out of scope for the CAT as an example of this. Also, digital assets if they failed the Howey test are non-securities. Covered liquidity staking receipt tokens, governance token with rights, and spot crypto ETPs are regulated under SEC. Spot crypto asset contracts are regulated under the CFTC. "Ancillary asset", such as fungible tokens sold via Simple Agreement for Future Tokens, used by crypto developers to raise capital from accredited investors before a token is live or functional, is proposed to use a hybrid regime (i.e. SEC as Primary Sales Regulator, and CFTC as Secondary Trading Regulator).[55]

The honorable goals of improving transparency and better identification of addressable liquidity are unfortunately no-match to this segmentations and money flows toward DeFi. TradFi establishments infuse trust into crypto ecosystem, while a toll gate to rent-seek from flows passing through their infrastructure is inevitable. The approach to "suck" new money from Digital Asset activities may be unconventional, but it is becoming the biggest force behind DeFi. See our October 2025 comment letter to the US Treasury regarding Innovative Methods to Detect Illicit Activity Involving Digital Assets.[56] We applaud the GENUIS Act to dollarize everything to turn the table against de-dollarization threats.

The CAT is wholly inadequate for the challenges of the twenty-first century. Past administrations did not set it up right in the first place. Make meaningful changes to right the course of the CAT's scope and design in addressing its fundamental flaws. Think big thoughts about CAT, or create new and better forensic tools, see Annex 2.

b. Mandatory Reporting Entities (Jurisdictional Scope)

Trade reporting is outdated. The frequent transmittal of data in-and-out and within the CAT system, unnecessary data-in-motion traffic is a costly waste and more susceptible to defects. See our response to Q1. CAT LLC could have directly aggregated high-speed, native pipelines to augment the missing order-level details that legacy systems like OATS and MIDAS lacked. Had regulators simply synthesized these existing proprietary feeds alongside centralized clearing and settlement systems (see part e), the SEC and SROs could have conducted comprehensive market-wide analysis without ever needing to construct a centralized, multi-billion-dollar CAT data vault.

[51] https://www.cboe.com/insights/posts/the-state-of-the-options-industry-q-1-2026/
[52] https://www.databoiler.com/index_htm_files/DataBoiler%20CFTC%20SEC%2020260430.pdf
[53] https://www.linkedin.com/pulse/modeling-trade-war-from-currencies-kelvin-to-zxsbe/
[54] https://www.databoiler.com/index_htm_files/DataBoiler%20Treasury%20Digital%20Assets%20202208.pdf
[55] https://www.databoiler.com/index_htm_files/DataBoiler%20CFTC%2020250818.pdf
[56] https://www.databoiler.com/index_htm_files/DataBoiler%20USDT%2020251017.pdf

According to SIFMA, the total private-sector compliance expenditure directly incurred by industry members to report data to the CAT has exceed $1.7 billion annually.[49] It is unjust for every single registered broker-dealer that receives, originates, handles, or executes orders in covered products to bear this cost when they have totally no control over CAT's budget, its operations, and receive absolutely no benefit from the CAT in return. Historically, the principle of *"no taxation without representation"* anchored the foundational fairness of American governance. In the context of current regulatory overreach, broker-dealers and end investors could justifiably assert a modern equivalent: *"no CAT-based fees without representation."*

Following the 2010 flash event, over 18,500 mini flash crashes have occurred in individual stocks, yet the **CAT did NOT generate a single alert**, nor did it in any way avert the trend of continuously rising Ultrafast Extreme Events. Broker-dealers who are doing things fairly and squarely should NOT be subjected to regulatory scrutiny and CAT cost burden. CAT should ONLY be used to capture suspicious activities from surveillance and regulatory market monitoring systems.

We recommend relieving the broker-dealers' burden. Their trade reporting obligations should be limited to the equivalent functions of OATS, COATS, EBS, and EDGAR filings (e.g. 13f, 13h), (and TRFs for fixed income if expanding the number of covered products). All major proprietary trading firms and broker-dealers subjected to Reg. SCI for coordinated cross-market Level 1 and Level 2 circuit breaker tests can and should be targets for scrutiny.

For SROs, they must provide their fastest connectivity and proprietary level 3 and historical market data to any revised systems or any new forensic tools as a requirement. If covered products are expanded per part a, then CFTC regulated DCMs, SEFs, and Derivatives Clearing Organizations (DCOs) should also be brought within scope. Any operators of TRFs, clearing and settlement systems, and the SEC operated EDGAR system should allow the new forensic tool to gather intelligence relevant to identified suspicious activities.

c. Covered Events (Lifecycle Scope)

Anyone who understands the trade lifecycle[57] would know the best source of trade data always resided in the clearing and settlement systems. The SEC and FINRA never lacked trade data in MIDAS and OATS in the first place. Explicitly excluding clearing data from the CAT is one of the fatal flaws (see part e).

It is understandable why the Commission staffers despise the process of submitting an Electronic Blue Sheets (EBS) request ONLY after they notice an anomaly on the tape, yet this procedure is a necessity to uphold privacy protection (see Q24). Casually asking if anyone has an interest (an RFQ for example) rather than submitting a "firmed quote" should NOT be subjected to regulatory scrutiny, hence is NOT within the scope of "Origination/ Receipt" or "Modification & Cancellation" – see our response to Q20 with regards to Requests for Quotes (RFQs). CAT operates as a continuous, mandatory daily ingestion loop that give regulators an instantly searchable repository without having to alert a firm that they are being investigated, which is a **civil liberties concern**,[20] see SECTION G.

Addressable liquidity is a fluid concept in Europe,[12] and attempts to embed FIX Market Model Typology (MMT) lifecycle flags into real-time market data feeds introduced several inefficiencies. Off-book on exchange OTC trades are subjected to deferral transparency regime for privacy protection of bilateral activities. SIP's pre-trade Consolidated Quotes System (CQS) and post-trade Consolidated Trades System (CTS) are substantially sub-par to self-aggregators' or existing market data vendors' non-display ultra (sub-microsecond) or low-latency feeds performance. All of these factors underscore that creating a centralized single source of truth within the CAT is unrealistic and unattainable. The problem lies within the unlimited desires – wanting CAT to be "everything everywhere all at once" to everyone.

[57] https://www.databoiler.com/index_htm_files/DataBoiler%20BigPic%20-%20CapitalMarket.pdf

Pleasing the SROs or anyone else to expand usage of CAT data is a departure from the Congress's original mandate and it may be considered as an alleged misallocation of CAT funding resources. CAT data covered up known issues on lit Exchanges (see our response to Q4). SROs have everything they need to produce an accurate, time-sequenced record tracking the entire lifecycle of an order from inception to finality if SROs are willing to synthesize their existing proprietary feeds alongside centralized clearing and settlement systems. The SEC, GAO, and DOJ should investigate if the CAT data was **weaponized for political or commercial reasons** that benefit the SROs and their vendors.

d. Mandatory Data Points (Granularity Scope)

For every covered event, the CAT NMS Plan requires collecting highly granular identifiers: **Unique Firm IDs**, **Unique Customer IDs** (CAIS tracking account holders and anyone holding trading discretion), and **Synchronized Millisecond Timestamps**. See our response to Q2 and below regarding our concerns with CAIS.

Using new jargon as approved by the Jan 2026 amendment, such as "Account Reference Data" (ARD) and "Customer Reference Data" (CRD) to supersede "Customer Account Information" (CAI) and "Customer Identifying Information" (CII) respectively does NOT make CAT less vulnerable to privacy and cybersecurity threats (lipstick on a pig is still a pig). The proposed retention of ARD and CRD together with the log-data that documents and reviews deletions of Customer names, address, years-of-birth, authorized trader names, SSNs/ ITINs under the Proposed Amendment is indeed allowing the possibility of reverse engineering to reconstruct the entire CAI and CII privacy information.

The procedure to summon sub-account and underlying beneficial owner information only upon an anomaly or suspicious manipulation activity being identified is a necessity to uphold privacy protection. The customer data of ordinary citizens should NEVER be routinely accessible to any government agency.[20] The reluctance of the Commission and CAT LLC operating committee to completely purge/ eradicate non-public and personally identifiable information (PII) exacerbates systemic privacy and security risks.

e. Explicit Exclusions (Out of Scope)

We can comprehend the rationale for excluding Primary Market Transactions. The CAT mandate is for Flash Crashes, NOT capital formation. There is no electronic "Order Lifecycle" to track IPOs. Separate regulatory regimes exist, where FINRA Rules 5130 and 5131 closely police IPO allocations and spinning practices. It was proper keeping IPO and primary issuances out of the CAT scope. However, per our response to part a, we strongly recommend bringing Futures and Swaps, as well as Clearing Data back within the scope of CAT (or preferably a new **forensic tool**) to meet the challenges of today and the future.

17. **Should the Commission modify CAT functionality** (e.g., lifecycle linkage, processing timelines, data retention)? What changes and why?

The whole concept of requiring everyone to submit data into a centralized vault itself is flawed and invasive. "Golden-sources of data" are prime targets attracting hackers to treasure hunt – i.e. honeypot risk. The **Millisecond Timestamp Tolerance** makes the CAT data useless. It is absolutely opaque for FINRA CAT LLC to rely on **closed-door contract** of nearly **$770,000 annually** to pay for a so-called "Market data vendor fees" (see Q4).

What actionable data elements or specialized processing does this "market data vendor" offer that could not be achieved more cost-effectively and with greater scalability through native AWS infrastructure?

How is it NOT a structurally redundant data loop if it is a basic CPU-bound data-wrapping process to regurgitate SIP and other public data mentioned earlier? How is it NOT introducing noise or distorting the NBBO sequencing of orders and trades in CAT when it serves no useful calibration or technical normalization purpose beyond fitting stale data?

SROs already generate comprehensive Level 3 depth-of-book proprietary feeds. Had regulators simply aggregated these existing proprietary feeds, there is no necessity for a CAT to set up cumbersome linkage modules to reconstruct the order sequence that is prone to errors and susceptible to noise.

We would like to reiterate our proposed approach per Appendix 2 of our May 2021 comment letter.[58] Analogous to the IRS' successful "my free tax initiative,"[59] this approach would allow the SEC and CAT participants to focus on those high-risk candidates for scrutinization. We envisage a crowd model to reduce unknown unknowns while enhancing security of the CAT. The benefits of our suggested approach are:

(a) dramatically reduce the CAT footprint or data storage and traffic by avoiding unnecessary redundant copies of data and minimize "data-in-motion";

(b) confine access of CAT data to 'targeted search' of relevant data that fits the "defined purposes"; and

(c) better intelligence for market monitoring by enabling and rewarding the crowd for identifying early warning signals to potential flash crashes or other trade irregularities.

MCP is a way to harmonize execution logic across fragmented venues. MCP exposes or abstracts venue-specific quirks – for example: "exchanges optimally restricted access to price information";[31] for another example: "how TRFs handle and preference orders can impact execution quality."[60] MCP sessions generate structured logs with metadata like tool provenance, invocation intent, and latency assumptions. Auditability and transparency of MCP could potentially replace SROs' proprietary reconciliation engines. It is ideally suited to meet the operational demands of regulators and compliance teams alike.

CAT puts an undue burden on broker-dealers, it is time to right course its scope and architecture rather than spending millions more trying to patch a broken centralized linkage model function. Those resources should be redirected toward a Real-Time Analytical Platform (RTAP) powered by Agentic AI (see our response to Q22).

18. **Should the Commission modify CCID generation** or the way customer-level analysis is performed? What changes and why?

Please see our response to Q22 for a holistic review of components of the CAT system linking the CAT-Order-ID to a customer identity via a two-layer mapping architecture (i.e., FFID and CCID).

Ditch the entire CAIS, it will save $21+ million a year. Upon identifying suspicious activities, then summon the account details from broker-dealers or through a vendor system on a case-by-case basis. CAT is allegedly an unauthorized "census" outside of SEC statutory authority (see our response to Q2). The Commission and SROs do not need to collect and store everyone's financial activity into a centralized vault to fulfill their Congress's CAT mandate on flash crash prevention. One does not need to know who is who in the entire universe but following a combination of publicly available digital trails, secret identities can be unveiled. See Annex 2 to marvel at the advancement of Agentic AI for better ways to detect, hunt down bad actors/ foreign adversaries and prevent flash crashes while respecting civil liberties, improving privacy and securities controls, at a lower cost and reduced footprint.

[58] https://www.databoiler.com/index_htm_files/DataBoiler%20SEC%20CAT%2020210503.pdf
[59] https://www.linkedin.com/pulse/hr-block-analogy-cat-combating-fraud-kelvin-to/
[60] https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3369025

Section E — Previous Changes to CAT Requirements

19. **Should the Commission modify requirements related to verbal activity on exchange floors?** What changes and why?

Commercial voice surveillance solutions (e.g. NICE Actimize) are widely used by firms for internal compliance, but the SEC continues to push for verbal activity integration into CAT anyway.[61] The decision or attempt to modify the CAT NMS Plan requirements rather than relying on existing commercial infrastructure stems from a deep division between **regulatory blockhead** and **operational market realities**.

(a) The Core Legal Issue: "Audit Trail Completeness"

Commercial voice surveillance platforms record audio, flag keywords, and archive transcripts for internal compliance officers. However, they do not output data formatted to automatically "stitch" an un-systematized verbal floor quote directly to a later electronic execution tape inside the CAT central repository. If a floor broker verbally announces a firm order or a market maker shouts a firm quote, that activity can alter market liquidity. The SEC maintains that without these verbal milestones formally ingested as structured event rows in the CAT database. NOTE: CAT data is NOT suitable for BestEx review (see our response to Q1, Q4, and Q22).

(b) The Private vs. Public Data Capture Chasm

While exchange groups own highly sophisticated surveillance tools, that data is siloed. For example, **i**f the SEC's Division of Enforcement wants to cross-examine an options trade sequence that began with a verbal shout on the CBOE floor and ended with an electronic execution on ARCA, they cannot easily stitch together two separate proprietary vendor databases (i.e. wanting CAT to be "everything everywhere all at once" to everyone).

(c) The Compromise: The June 2025 "Verbal Quotes" Amendment

The exchange participants aggressively fought back against the initial rule, arguing that forcing floor brokers to manually log every spoken word into a CAT-compatible electronic schema would completely paralyze the open-outcry trading process and cost billions. This pushback forced the SEC into a major structural compromise in its June 2025 Order.[62] **The 2030 Sunset Extension** permanently excluded unstructured "upstairs" telephone/chat logs from CAT reporting. For verbal floor activity, the SEC capitulated to the industry's cost concerns by deferring the mandatory electronic reporting requirement **until July 31, 2030**. In essence, the SEC is hoping that by 2030, the very commercial AI voice-to-text vendors will evolve to the point where they can automatically translate spoken floor chaos into perfectly formatted, structured CAT data sheets without human intervention.

Automated Market Makers (AMMs) may widen spreads to avoid "pick-off" rather than continuously provide tight two-sided quotes. When volatility spikes or an "asymmetric information shock" hits, an algorithm's code may prioritize protecting its own balance sheet first. Do we NOT trust the human floor agents any less than some of these AMM algorithms to maintain a continuous orderly function of markets? Humans are subjected to strict compliance oversight, e.g. the taping mandate, no open lines, device auditing and registration, pre-vocalization/ electronic systematization rule, etc.[63] Exchanges completely isolate the physical floor environment from unauthorized electronic data transmissions with total video bans and network jamming prevention. Do not over-regulate human floor agents while remaining blindsided by algorithmic privileges (see Q21).

[61] https://www.sec.gov/files/rules/sro/nms/2024/34-101648.pdf

[62] https://www.govinfo.gov/content/pkg/FR-2025-06-20/pdf/2025-11331.pdf

[63] https://www.niceactimize.com/financial-markets-compliance/communication-compliance/communications-voice-management ; https://www.govinfo.gov/content/pkg/FR-2017-05-23/pdf/2017-10588.pdf ; https://www.govinfo.gov/content/pkg/FR-2021-05-05/pdf/2021-09443.pdf ; https://www.sec.gov/files/rules/sro/nms/2024/34-100727.pdf

20. **Should the Commission modify requirements related to electronic requests for quotes?** What changes and why?
In general, RFQ is a non-binding expression of interest—an inquiry to see who has liquidity. We can comprehend concerns about "Phantom Orders" where signals are extracted using electronic RFQs to scan the market but lack trade executions (see Q24). There is a distinction between "immediately actionable responses" (e.g. a market maker responds to an exchange RFQ with an electronic firmed quote (e.g., via standard FIX protocol) and "Non-Immediately Actionable Responses" (the SEC granted permanent exemptive relief in legitimizing a responder to send an indicative price or an incomplete quote that requires a secondary "pass" or final electronic confirmation before a trade can lock).

Sending out an RFQ or streaming market data technically costs little to almost nothing. Yet, one ought to pay for royalties in streaming copyrighted materials published by others. There is tremendous value in composing trades and publishing quote and trade algorithms. Allowing it to stream freely at no cost is like a pirate copy of an MP3 song. By transforming quotes and trades into music, with the appropriate obfuscation to preserve confidentiality of strategies, the rights to claim ownership of the data by broker-dealers can be asserted.[64]

The matter here is about capabilities difference between lit Exchanges and Wholesalers, ATSs, and SDPs. There may not be good or bad guys, but a continuous race to unveil versus concealment of trade intents to optimize performance, reduce slippage, and maximize profit. Regulatory enforcements ought to be objective. Subjectively **weaponizing CAT for political or commercial reasons** that benefit the SROs is a departure from Congress's original mandate (see Q1).

The SEC may update the relevant RFQ or block trading requirements if the new policy can improve price discovery and capital formation. However, the Commission and SROs cannot weaponize CAT for political or commercial reasons. Forcing firms to log every electronic inquiry including "Non-Immediately Actionable Responses" is an intrusion upon privacy.[65] CAT exacerbate vulnerability where a breach (see our response to Q28) would give hackers a blueprint of every major institution's private, unexecuted trading strategies.

Again, CAT in its current form is incompatible with Vice President JD Vance's remarks about "American AI will not be co-opted into a tool for authoritarian censorship." We concur with this blog post by the CATO Institute with regards to the SEC's push to drag Electronic RFQs into the CAT that implicates the Fourth and Fifth Amendment Rights of Investors.[66] Uphold citizen data sovereignty – neither massive tech monopolies nor unchecked federal agencies should have the right to scrape, store, and profile private American activity.[20]

21. **Should the Commission modify requirements related to port-level settings?** What changes and why?
Institutional trading firms do not always specify handling instructions on an order-by-order basis. Instead, they configure the port itself to automatically apply default behaviors to every incoming order—such as price sliding, short-sale presentation rules, or self-trade prevention. We do recognize that in algorithmic driven capital markets,[67] when the code, configuration, or speed threshold of a broker's outbound port mismatches the exchange's inbound data port, catastrophic, ultrafast market events can occur, see below for historical crash event examples:

- Single-Sided Inbound Buffer Loop in 2010 Flash Crash: AMMs started absorbing massive sell volume, the raw data throughput entering the exchange gateway ports skyrocketed exponentially. Several major market-making broker desks had their internal communication ports bottleneck. Because their outbound confirmation loops could not

[64] https://www.databoiler.com/index_htm_files/DataBoiler%20Noumenon%20Equity%20Market%20Structure.pdf
[65] https://www.sec.gov/newsroom/speeches-statements/peirce-nms-cat-2020-08-21
[66] https://www.cato.org/blog/secs-market-surveillance-system-implicates-constitutional-rights-investors
[67] https://www.linkedin.com/pulse/from-latency-ai-algo-driven-capital-markets-kelvin-to-xu5te/

process data as fast as the exchange ports were broadcasting fills, a massive "latency queue" formed. The algorithms, suddenly operating on delayed data, assumed they were flying blind and triggered their automated emergency kill-switches to pull out quotes entirely. In turn, the market was in a complete liquidity vacuum that caused prices to plunge to pennies in milliseconds.

- The Knight Capital Crash (2012): When live order flow hit that eighth server where updated code had failed to deploy, its legacy code repurposed a dormant port-flag. Instead of sending standard, metered orders, it trapped the incoming client data into an infinite routing loop. The mismatched port blasted millions of un-throttled, unintended market orders directly into the exchange, instantly driving dozens of NMS stocks into localized mini-flash crashes before the server could be manually disconnected.
- FIX Protocol Rejection Loops and Mini Crashes: Exchanges enforce strict "Message-per-Second" (MPS) throttling limits on every physical 10Gb or 40Gb cross-connect port a broker rents. If a broker's algorithm miscalculates and floods the port past its hard-coded ceiling, the exchange engine automatically rejects the excess messages. If the broker's system is misconfigured, it may interpret those exchange rejections as a technical connectivity loss and instantly attempt to cancel and resubmit all resting quotes over that same port. This creates an exponential, self-reinforcing loop. Within microseconds, the broker's port completely jams, preventing their market-making system from updating quotes while the broader market moves. HFTs' arbitrage algorithms detect this stale, locked liquidity and immediately "pick off" the frozen quotes, causing an instant, sharp price dislocation in that stock.

Do not over-regulate human floor agents (see Q19) while remaining blind to algorithmic privileges. We have reservations with Dual-sided reporting. In practice, routing firms do not maintain the receiving firm's port settings in their own internal books and records.[68] Exchanges profit heavily from proprietary connectivity and co-location infrastructure, yet the administrative burden and costs of reconciling those custom environments are being externalized onto CAT and everyone.

(a) The Co-Location Profit vs. CAT Expense Loop

Exchanges treat co-location and customized high-speed ports as high-margin, proprietary product suites. They create unique, non-standardized port configurations (like *Self-Trade-Block-01*) to "**LOCK IN"** clients and "optimally restricted access to price information."[31] **DOJ should investigate these anti-competitive practices**. **The Cost Shift**: Rather than building their own native compliance interfaces to resolve the naming conflicts their proprietary systems create, exchanges push the raw data into the CAT. **The Subsidy**: Because the CAT infrastructure is a centralized utility, the massive computing bills required to run complex string-matching algorithms over billions of rows are spread across the entire industry.

(b) The CAT Funding Battle (The Fee Dispute)

The Unfair Burden: Broker-dealers are being financially penalized by the micro-technical push by the SEC to ingest port-level settings into CAT. It forces them to build and maintain the expensive data-formatting loop (clean, wrap, and dump billions of raw data rows into a central cloud vault) – an enormous, continuous, industry-wide data-sharing and pre-linkage process. This data pipeline is required to decode data that the for-profit Exchanges could

[68] https://fif.com/index.php/working-groups/category/271-comment-letters?download=2905:fif-exemptive-request-letter-to-the-sec-on-port-settings&view=category ; https://www.sec.gov/comments/4-853/4853-618547-1815754.pdf

easily standardize or flag themselves natively at the point of origin. Dual-sided reporting can never be cost justified, see Q16 part c and Q22. The SEC should make the **Jan 2026 Exemptive Relief Order**[69] permanent.

We recommend that the SEC revise the port level requirements to **hold SROs accountable**, **eliminate anti-competitive practices**, and abandon the legacy SQL-based querying over structured data, which currently necessitates the daily ingestion of every port configuration change into the CAT's centralized AWS cloud repository. We counter suggest a **parallel process** to streamline the back-and-forth *"unlinked/ Error"* communications between SROs and Broker-Dealers, and making related resolutions transparent to the SEC, see below steps:

- Step 1: During the trading session, a Broker-Dealer routes an order to an SRO's co-location engine. The broker does NOT clean, wrap, or format this data into a complex CAT-compliant schema. They simply log the port setting string naturally within their own secure boundary layer as STP_DEFAULT. Concurrently, the exchange logs it natively as Self-Trade-Block-01.
- **Step 2:** Instead of waiting for an overnight batch or a structured data dump from the broker, a mirrored copy of the raw session handshake is fed straight into a Localized AI Compliance Agent running at the Exchange edge. The SRO's AI agent uses semantic parsing to compare the two raw strings instantly. Within microseconds of execution, the agent determines that STP_DEFAULT and Self-Trade-Block-01 represent an unmapped mismatch.
- Step 3: The SRO's AI agent dynamically wraps this mismatch into a lightweight, schema-agnostic Mismatch Notification Document. The agent instantly dispatches this document via a secure API pipeline to (A) sent straight to the Broker-Dealer's incoming operations queue, and (B) simultaneously sent directly to the SEC – a new schema-agnostic database (we recommend using **Progress MarkLogic**, see Annex 2).
- Step 4: Because **MarkLogic** stores data as-is, maintains a universal index, and includes a native triple store, it captures the relationship metadata (SRO_X → Flagged_Mismatch → Broker_Y on Port_Z) without the rigid schema design and heavy Extract-Transform-Load (ETL) that a traditional relational database would require. The mismatch is written once and indexed automatically. This creates a live, open "Dispute/Mismatch Thread" in the SEC's database within seconds of the actual trade.
- Step 5: Upon receiving the notice on T+0 at the broker-dealers' secure boundary layer, the firm can choose between: (i) an autonomous AI compliance agent to immediately query internal logs and resolve the mismatch, or (ii) routing the notification in a secure compliance queue to a human compliance analyst or standard automated workflow to be investigated and addressed manually within the allowable regulatory duration.
- Step 6: regardless of option (i) or (ii) the broker-dealer chooses, the broker-dealer provides a standardized resolution document where they push a patch code to **fix** and align their data string with the exchange, or reject the notification and logs a **Formal Rebuke Document** attaching technical system logs to prove the exchange's port logic is what actually misbehaved.
- Step 7: the broker's response (**Fix** or **Rebuke**) is simultaneously copied to the SRO and the SEC's sematic database. If it was a **fix**, the notified case is marked as *"resolved"*. If a **rebuke** is filed, the SEC tracks the thread's lifespan. If the allowable regulatory duration (T+3) is expired without a commonly agreed resolution between the SRO and Broker-dealer, the system would generate an alert prompting the SEC to step-in.

[69] https://www.federalregister.gov/documents/2026/01/27/2026-01611/order-granting-exemptive-relief-pursuant-to-section-36a1-of-the-securities-exchange-act-of-1934

This decentralized AI-driven architecture eliminates formatting burden on broker-dealers. The compliance focus returns to the for-profit exchanges where it belongs. By leveraging real-time edge AI alongside the semantic flexibility of **Progress MarkLogic** database, the SEC gains immediate visibility into market-wide operational friction without managing petabytes of raw data loops. Furthermore, by offering firms a choice between autonomous AI responses or traditional manual workflows, this model respects varying operational maturities while creating an immutable, transparent audit trail that eradicates legacy regulatory latency.

22. **Should the Commission modify requirements related to representative order linkage?** What changes and why?

CAT relies on an internal system generated identifier called the **CAT-Order-ID** to reconstruct the order sequence:

- The linkage module looks at a single broker-dealer's internal logs; matches that firm's localized "orderID" across the life cycle of the order—tying the Origination/Receipt to any subsequent Modification, Cancellation, or Execution.
- When Broker A routes an order to Broker B, the module uses a composite Route Linkage Key to link the two firms. This key is built by combining fields like the routedOrderID, the sender/receiver Industry Member IDs (IMID), the specific trade date, and the ticker symbol. If Broker A's "Sent" record matches Broker B's "Received" record across these data points, the engine pairs them accordingly. **Dual-sided reporting can never be cost justified.**
- When a broker routes an order to a public stock or options exchange, the repository matches the broker's routing data with the exchange's native electronic order book logs.

The system links the CAT-Order-ID to a customer identity via a two-layer mapping architecture:

- The Local Layer: The broker-dealer tags the transaction record with a unique, persistent Firm Designated ID (FDID). The FDID typically maps back to the account relationship. Crucially, the broker is strictly prohibited from using actual account numbers, social security numbers, or names in this field to prevent identity theft.
- The Central Layer: In the separate CAIS database, the firm securely uploads the "true" customer data mapped against that specific FDID. The central system runs an algorithm to generate a universal CAT Customer ID (CCID).

Clock drift, or a typo in a broker's routedOrderID field can cause a **Linkage Error**. There are three main causes of Data Quality problems:

- Self-serving human-induced biases and institutional favoritism: e.g., crappy data (i.e., 100,000 messages at any given point in time that mussed everything up) and/or hold off the advancement of the Consolidated Tape for ecosystem degradation to exacerbate the gap between Proprietary Products and Consolidated Tape. Consolidated Tape with a wide bid-ask spread and delay refreshing of the NBBO, i.e., makes the tape unsuitable for Best Execution (BestEx) analysis, in turn, majority of market participants are not aware that they have been receiving inferior prices.
- Poor controls, system glitch, operational resiliency: e.g., technical issues with a trading venue's volatility interruption mechanisms can send shares on a wide ride (e.g. Berkshire price erroneously down 90+% in June 2024);[70] system glitch such as the Knight Capital $440 million software error;[71] and data center redundancy failure[72] (outage),[73] operational resilience[74] (some see the price some do not), and cybersecurity issues.[75] Non-harmonized logics among SROs caused back-to-back halts.[35] SROs and all major proprietary trading firms and broker-dealers under Reg SCI are accountable

[70] https://www.nytimes.com/2024/06/03/business/nyse-trading-glitch-berkshire-hathaway.html
[71] https://www.henricodolfing.com/2019/06/project-failure-case-study-knight-capital.html
[72] https://www.nasdaqtrader.com/content/newsalerts/2017/eta/SIP_Contingency.pdf
[73] https://www.tradersmagazine.com/featured_articles/brokerage-outages-highlight-vulnerabilities-in-the-infrastructure-of-online-trading-platforms/
[74] https://www.whitecase.com/insight-our-thinking/financial-regulatory-observer-2022-operational-resilience-uk-eu-and-us
[75] https://www.nist.gov/cybersecurity/international-cybersecurity-and-privacy-resources

to adopt certain volatility interruption best practices[76] to ensure price continuity. Reluctant to conduct data consistency checks with peer TRFs is attributed to incomplete or potentially erroneous information.

- Inadvertent or honest mistakes during trade reporting, such as duplicated positions and/or misuse of taxonomy flag(s) by investment firms during trade reporting to TRFs/ SIP. Broker-dealers already face undue burden (total private-sector compliance expenditure directly incurred to report data to CAT has exceeded $1.7 billion annually). Putting ALL THE BLAME on investment firms for the failure to generate the CAT-Order-ID is WRONG.

Fragmented data did NOT and will NOT leave regulators stranded. The Commission and the SROs refusal to address initial bias and overhaul the fundamental architectural flaws[9] of CAT (a giant centralized data vault with wasteful data at-rest and unnecessary data-in-motion)[11] did. False belief in lifecycle flags to push for a Unique Transaction ID further burdening broker-dealers to implement two-side reporting, which exacerbate CAT's unsustainability. CAT operating committee allocated most of CAT's annual operating budgets of $188 million to $272 million to "Cloud hosting services" but leave little money in capitalizable technology development and almost nothing in data analytics. How could the essential layers of intelligence ever be built like this?!

Reference to our response to Q10, we believe Dr. Gregg Berman (a highly respected CAT advisory committee member, former Associate Director at the SEC that led the Thesys built MIDAS project and is currently a Director of Market Analytics and Regulatory Structure at Citadel Securities) supported the use of a daisy chain for CAT. For over half a decade, substantial CAT funding was poured into engineering rigid tracking loops for representative orders (where a broker-dealer bundles multiple customer orders into a single "representative" block trade to execute in the market). The sudden pivot to modify these requirements is a tacit admission that the original "daisy chain" approach was a flawed, incredibly expensive dead end.

The CAT operating committee could and should have known that representative orders inherently break traditional daisy chains because a broker is constantly slicing, dicing, combining, and allocating shares across different proprietary desks and customer accounts, a simple linear chain cannot structurally account for the data relationships. This "failed experiment" resulted in massive volumes of unlinked representative orders, triggering endless data validation errors. Compliance teams were forced to spend thousands of man-hours daily manually fixing broken links that the rigid system could not automatically parse.

This is a huge regulatory tech-stack waste. Broker-dealers, clearing houses, and third-party connectivity vendors (like Broadridge) had to completely rewrite their core execution logic and order/ execution/ portfolio management systems (OMS/ EMS/ PMS) to feed data into a rigid daisy-chain repository. Much of that engineering work has effectively been rendered obsolete. We are suspicious that the SEC and SROs attempt to modify requirements toward a looser, more flexible structural association would yield any fruitful result. Why should the industry shoulder another round of R&D experiments run by the same CAT operating committee?

Per our response to Question 425 in our July 2016 comment letter,[77] "Daisy Chain is more favorable than a unique ID approach. Daisy Chain can be configured by connecting each component to another similar component, rather than require multi-interfaces to directly connect the system's core to each of the components being used. Although daily chain is easy and cheap, it has a significant drawback – Daisy chain is suitable only for fixed functional units that

[76] https://www.xetra.com/xetra-en/trading/protective-mechanisms/protective-mechanisms-in-continuous-trading + participate in at least one coordinated cross-market Level 1 and Level 2 circuit breaker test per year
[77] https://www.databoiler.com/index_htm_files/DataBoiler%20CAT613%20Comments.pdf

cannot be removed… 'Tree' and 'Mesh' are more effective ways than Daisy Chain to cope with modern topology design requirements… we can help optimize between the cost of control and the most desirable 'link' of orders."

The constant modification of CAT rules – from RFQs to port settings, and now representative order linkages – proves that centralized data hoarding cannot keep pace with dynamic modern markets. The CAT system is structurally stuck trying to solve 2026 and beyond problems with 2012 and older database technology. Rather than spending millions more trying to patch a broken centralized linkage model, redeploying the resources to Agentic AI.

Instead of forcing a thousand different broker systems to structurally alter how they link data before sending it to an AWS vault, an Agentic AI system could analyze the data directly at its source:

(a) Contextual Reconstruction vs. Rigid Linking

An Agentic AI auditor does not need a pre-formatted, hand-labeled daisy chain. By deploying autonomous agents directly at the exchange and broker source levels, the AI can perform **probabilistic and temporal reconstruction**. It can look at a block execution on an exchange tape, instantly scan local broker order books, and mathematically deduce—via execution time, size, and routing patterns—exactly which representative order matches which suspicious account.

(b) Sub-Second Real-Time Threat Analysis

The current CAT system is inherently retrospective; data is dumped at End-of-Day and repaired over a T+3 window, meaning regulators are always looking at the market through a rearview mirror. An agentic RTAP framework sits directly on the live data pipelines. If a manipulative cross-market loop occurs, local agents flag the anomaly as it happens, pulling only the relevant context for SEC review.

(c) Bypassing the Data-Wrapping Tax

Implementing an agentic framework would completely eliminate the need for third-party commercial data-wrapping loops. Because LLM-driven agents possess native code-generation and semantic parsing capabilities, they can query data across different legacy firm databases without requiring the firms to pay vendors to standardize the schemas into a centralized repository.

See Annex 2 to marvel at the advancement of Agentic AI for better ways to detect, hunt down bad actors/ foreign adversaries and prevent flash crashes.

✉ info@databoiler.com https://www.databoiler.com

Section F — Potential Changes to Other Data Sources

23. **Should the Commission retire partially duplicative systems** (e.g., OATS, COATS, EBS)? Which ones and why?

Under normal circumstances, newer tech-stack replaces their older versions if its resiliency is proven to be stable to deliver superior performance at lower cost. However, following the Eleventh Circuit Court vacating the 2023 Funding Order, the subsequent lawsuits by Citadel Securities to freeze reserves, and a projected "coffer depletion date", the CAT faces an existential structural crisis. CAT was given 10+ years to experiment, not once (Thesys) but twice (FINRA CAT LLC), to do a $2+ billion proof-of-concept about a "centralized single source of truth" that is doomed to failure.

If the SEC sunsets the CAT and attempts to resurrect the legacy decentralized predecessors – OATS (Order Audit Trail System for equities), COATS (Consolidated Options Audit Trail System), and EBS (Electronic Blue Sheets for customer mapping) – the financial industry would face a massive engineering paradox.

Reference to our response to Q2, Q16 part d, Q18, and Q22, if the SEC however strips out the centralized linkage and the CAIS customer database from the CAT system and reengineers the remainder, the engineering scope shrinks drastically. The following are our estimates:

(a) Hardening OATS/COATS Batch Uploads (Zero-Trust)

The Vulnerability: Legacy OATS relied on basic FTP/SFTP file drops at the end of the day. By modern standards, these static entry points are highly vulnerable to credential-stuffing and data-interception.

The CapEx: $100 million to $150 million (Industry-Wide) depreciable over 7-10 years

The Upgrade: Every broker-dealer would replace legacy connections with **Zero-Trust Network Access (ZTNA)** and mandatory Multi-Factor Authentication (MFA). Data in transit would be forced onto TLS 1.3 encrypted tunnels with automated API key rotation to eliminate static password risks.

(b) Upgrading EBS with At-Rest Cryptography

The Vulnerability: The Electronic Blue Sheet (EBS) system transmits clear text or lightly masked subscriber identity data directly to regulators upon request. A breach of an EBS storage silo exposes raw investor profiles.

The CapEx: $100 million to $180 million depreciable over 7-10 years

The Upgrade: Firms would keep customer data strictly localized within their internal firewalls but update the database infrastructure to support **AES-256 at-rest encryption** and field-level masking. Regulatory queries would pull data through a decentralized secure enclave, ensuring that clear-text PII is never exposed on an active network segment.

(c) Continuous Penetration and SIEM Auditing

The Vulnerability: Legacy audit trails were "set and forget" systems, lacking active monitoring to detect when an adversary was scanning the ports or trying to exfiltrate trade logs.

The CapEx: $50 million to $120 million depreciable over 7-10 years

The Upgrade: firms would integrate their regulatory reporting logs into centralized **Security Information and Event Management (SIEM)** engines. This ensures real-time anomaly detection, automated system patching, and continuous penetration testing to defend the decentralized nodes against zero-day exploits.

The money required to patch legacy systems is roughly what the industry is fighting over right now regarding the CAT. The CAT's current architecture is under intense scrutiny because its privacy and security protocols are deemed

inadequate by the industry. The SEC has already been forced to spend millions and make dramatic changes because the CAT's security setup failed to meet modern standards.

To stop massive lawsuits from Wall Street trade groups, the SEC issued an Exemptive Order and subsequent amendments to permanently eliminate the reporting of raw PII into CAT. Instead of completely purging customer tracking to save money, they had to design a new system to generate anonymized CCIDs using external reference data nodes. This change reduced potential cost-saving measures by millions just to establish a safer data structure.

The legacy perimeter model that CAT relies on is highly vulnerable to the automated, distributed query data leakage (see Q27). The CAT repository processes billions of algorithmic data messages every day. Securing a database of that size against modern cyber threats requires continuous, expensive upgrades. The CAT needs constant patches for its cloud-based secure enclaves, maintain strict field-level masking, and run continuous automated penetration tests – i.e., a **money pit**.

- $80-120 Million in near-term development Capital Expenditure, plus an ongoing annual operational overhead of $40+ Million. This funding is required to transition the central database into quantum-resistant encryption, maintain massive multi-tenant secure cloud enclaves, and administer continuous AI-driven penetration testing.
- **Caveats:** This security fix does NOT avert the risk of CAT being a prime target for sophisticated cyber warfare. As transaction volumes grow exponentially, the cost to store, link, and continuously scan a centralized repository scales linearly, creating a permanent funding bottleneck.

It is unrealistic to fulfill unlimited desires – wanting CAT to be "everything everywhere all at once" to everyone. **The CAT is broken beyond repair** and is wholly inadequate for the challenges of the twenty-first century (see Q3). The ONLY way to reduce the CAT's footprint is right coursing its scope and architecture. Instead of choosing between keeping OATS/ COATS/ EBS or CAT, the Commission should make bold move to *shift away from continuous trade reporting* (in CAT or these legacy systems). Instead of forcing market participants to constantly transmit billions of raw, duplicate logs to a central *honeypot*, the data would remain locked at its native origin: inside an exchange's matching engine, a clearing house's ledger (like the DTCC or NSCC), or an SRO's internal database. Regulators would instead deploy automated*, federated analysis models* directly to those environments (see Q27 and Annex 2).

- **Total Elimination of Systemic Reporting Bloat:** Today's CAT budget is strained because it forces thousands of broker-dealers to clean, format, and re-transmit data that already exists inside the infrastructure of the clearing houses and exchanges. Analyzing data directly at the source completely eliminates this data duplication, instantly saving the industry billions in redundant transmission pipeline overhead.[78]
- **The Ultimate Cybersecurity Framework:** Because data never leaves its native environment, the massive cyber-target of a centralized database vanishes entirely. Even if a bad actor manages to compromise a specific regulatory query tool, they gain zero access to raw, centralized pools of market data or investor information.
- **100% Accurate Transaction Timestamps:** Today's audit trails frequently suffer from linkage errors because of clock drifts and other issues (see Q4 and Q22). Analyzing data directly at the source—such as a clearing house or an SRO matching engine—ensures regulators view transactions using a single, immutable master clock, completely eliminating sequencing disputes.

[78] https://natlawreview.com/article/cats-ninth-life-secs-sweeping-review-could-fundamentally-reshape-consolidated-audit

- **Natural Regulatory Safeguards:** Regulators only access data on a strict "need-to-know" basis. The analytical tools only pull the specific transaction trails required for a live investigation, preserving foundational civil liberties and protecting proprietary trading strategies from passive surveillance.

Regulatory market monitoring ought to be modernized alongside market and technological advancements, such as Agentic Overlay. The goal is to assess changes in market dynamics, understand where frictions and liquidity concentration may shift at rapid pace or irregularities and dysfunctional market behaviors occur. Analyze the stress points and weak links where the next **flash crash** or systemic risks may emerge, and develop mitigation protocols with more effective and timely volatility interruption mechanisms.

24. **Should the Commission modify or replace the EBS system?** What changes and why?

It is easy to criticize the legacy EBS system if compared to modern technologies. Those undermining the merits of EBS should learn to appreciate that certain critical infrastructures and systems developed during, or before, that era were built with comprehensive consideration of various factors, such as privacy of customer account information. Back then, developers knew the best place where information should reside, to be the most accurate (e.g. clearing and settlement). Their designs were often purpose-driven to achieve specific goals – EBS is highly effective for targeted, rearview-mirror insider trading investigations. Hence, it deliberately uses ONLY the essential core data and emphasizes practical components to achieve optimal performance and realistic development budgets.

We are aware that if an institutional asset manager executed a massive trade through an omnibus account, or an executing broker, the EBS report would often only display the name of the clearing firm or the high-level institutional entity. To find the actual human or sub-account behind the trade, the SEC had to issue a second or third round of EBS requests down the clearing chain. It is understandable why the Commission staffers despise the process of submitting an EBS request ONLY after they notice an anomaly on the tape, yet this procedure is a necessity to uphold privacy protection. Moving PII from one place to another is NEVER a good idea, see Q2, Q16 part c, Q18, Q22, Q26, and Q27).

We can comprehend concerns about "Phantom Orders" where a firm sent 10 million electronic RFQs or algorithmic quotes to move a targeted securities price but only executed one trade. EBS may have a blind spot given it does not capture cancellations, modifications, or routing steps. Before drawing conclusion about alleged market manipulation, there are objective steps to distinguish "firmed quotes" from non-commitments. If the Commission wants to tighten rules around RFQs, and block trading requirements, then let's have separate rulemaking and assess whether it would improve price discovery and capital formation or not. formation. However, the Commission and SROs cannot **weaponize CAT for political or commercial reasons**. Forcing firms to log every electronic inquiry including "Non-Immediately Actionable Responses" is an intrusion upon privacy, see our response to Q20.

CAT should NOT be built solely for the convenience of regulators who remain wedded to legacy SQL-query frameworks within a centralized data vault. It is unrealistic to fulfill unlimited desires – wanting CAT to be "everything everywhere all at once" to everyone. CAT CAIS is a super expensive and intrusive experiment, putting undue burden on broker-dealers to map out the Ultimate Beneficial Owner (UBO) and unique Large Trader IDs (LTID) beforehand, see Q25. There is no need to reinvent the wheel, the Commission could have leveraged the bespoke model already built by vendors such as my former employer Broadridge that have API links to most if not all US broker-dealers for beneficial shareholders' information.

The SEC and SROs already possess essential information for flash crash prevention though they are from dispersed systems (see Q16). Annex 2 showcases how AI agents can be deployed when the SEC and SROs suspect anomaly and issue a federated encrypted query across the network of dispersed systems to expedite regulatory scrutiny. Should

any information need to be obtained with higher urgency or additional contents are required, then SEC must go through separate rulemaking. Conveniently pointing to the CAT to continually expand its scope to satisfy an unlimited desire to obtain anything the SEC and SROs want, while using Section 31 fees to impose financial transaction tax on everyone is unjust. The CAT has deviated from Congress's original mandate; it is an unauthorized "**census**" that falls outside the SEC's statutory authority, see Q2 and Q26.

25. **Should the Commission modify LTID requirements** or the large trader reporting system? What changes and why?

The SEC implemented the Large Trader Reporting System following the 2010 Flash Crash to efficiently identify market participants conducting substantial trading activity. 13h filing to the SEC is required if a person or institution becomes a "Large Trader"[79] and exercise investment discretion over accounts that trade NMS securities reaching either "2 million shares or $20 million in a single calendar day" or "20 million shares or $200 million in a single calendar month."

Historically, broker-dealers only provided Large Trader ID (LTID) records when the SEC manually requested Blue Sheets. Under the CAT requirements, broker-dealers must automatically stamp the LTID onto every single CAT Reportable Event (routes, quotes, cancellations, executions) on a T+1 basis. If an institution triggers the volume threshold but fails to self-identify via Form 13H EDGAR filing, the clearing broker is legally required to flag them under an Unidentified Large Trader ID (ULTID) inside the CAT repository, essentially reporting them to the SEC for compliance failure.

When CAT as a centralized repository reconstructs market sequences using third-party commercial data feeds – such as the data provided to the plan processor by Algoseek – to reconcile the SIP, it is impossible not to have sub-millisecond sequence drift (see footnote 10 about initial bias and Q4). Because LTIDs are automatically stamped onto these drifting sequences, the **CAT can be weaponized to penalize Wholesalers** and any "Large Trader" for artificial anomalies created by the SRO's own database.

To illustrate, if an HFT cancels an order at 10:00:00.000001 and an exchange logs a new quote at 10:00:00.000002, the order should be dead before the quote prints. However, if the consolidated SIP data processing loop buffers those messages incorrectly, CAT might sequence them in reverse. A mis-sequenced string of rapid-fire quote adjustments can look like deliberate spoofing, phantom liquidity layering, or quote stuffing. The automated system flags the event as an anomaly and attributes it directly to the firm's stamped LTID, shifting the burden of proof to the market participant to prove their internal logs were the accurate ones.

For another example, **ATSs Dark Pools** and **SDPs** depend heavily on customized execution logic, price sliding, and internalization parameters configured at the port level. If an ATS matches a block trade internally and reports it to a TRF, that trade must be synchronized with public exchange data. If the processing software misaligns the TRF print against a stale public quote due to data-ingestion lag, it creates the illusion of a trade-through violation[6] or a failure of Best Execution. If the sequence is structurally distorted, an SDP providing genuine liquidity during a period of high volatility might be incorrectly classified as "predatory" or "toxic", harming the firm's reputation and attracting meritless enforcement scrutiny.

Weaponization of CAT also arises from **enforcement asymmetry**. If a broker-dealer submits a sequence with a microscopic timestamp error, they face immediate FINRA CAT compliance fines. If architectural cloud-buffering loops

[79] https://www.finra.org/rules-guidance/guidance/reports/2021-finras-examination-and-risk-monitoring-program/large-trader-reporting ; https://www.catnmsplan.com/faq/q37

within the regulatory pipeline cause a sequencing error in the CAT transaction logs, the discrepancy is typically treated as a systemic artifact. The LTID holder is then left with burden for legal defense to disprove the data mismatch.

The SEC, GAO, and DOJ must investigate if CAT data was **weaponized for political or commercial reasons** rather than serving the Congress's mandate. Public confidence in market integrity erodes as the SEC continues to tolerate chronic national market disorders and data distortions within CAT and SIP. This instability creates a chilling effect discouraging market participations and places undue burden on market participants. In turn, capital will migrate toward digital assets, gaming platforms, and offshore jurisdictions, jeopardizing the dominant position of the U.S. financial system.

Section G — Civil Liberties and Privacy Considerations

26. **Are there additional privacy or civil liberties concerns** related to CAT or other audit trails? What changes should be made to address them?

The SEC **mischaracterized** EBS (see Q24) to downplay CAT's privacy and civil liberties vulnerabilities. The SEC concept release pointing to previous iterations of the audit trail – including the EBS for "response to a single EBS request could contain tens or hundreds of thousands of plaintext SSNs… transmission of SSNs and account numbers in plaintext risks the unauthorized disclosure of personal data… are not aware of any EBS controls to require the encrypted storage of SSNs and other PII." This ignores the fundamental operational difference. While legacy EBS file formats lacked modern encryption, the data was never aggregated systematically in one place. The SEC frames a fixable software formatting issue (encryption-in-transit) as a reason to favor the CAT, completely blinding itself to the fact that a centralized CAT repository creates an infinitely more dangerous, all-inclusive "honeypot" target.

The Commission argues that "because the EBS system routinely collects account numbers and underlying customer identities to help SROs investigate market anomalies, it presents similar systemic privacy risks to the CAT" completely misrepresents how the Fourth Amendment applies to the two systems. Under the EBS system, **investor data is only transmitted after a formal regulatory request or summons** targeting a specific, pre-identified trading anomaly. The CAT, by contrast, enforces a dragnet approach—**systemically vacuuming up every market participant's daily investment history** regardless of suspicion. Blockhead's baseline understanding of privacy is stuck in classic 20th-century paradigms (like protecting PII and raw database access), they are completely blind to algorithmic data leakage (e.g. Membership Inference, Reconstruction, and Model Extraction Attacks, see SECTION H). The SEC **conflating** "On-Demand" with "Systemic Ingestion" is completely **WRONG** and **misguiding the public**.

The Commission's attempt to create a Regulatory Safe Harbor by asking if there are additional or unaddressed privacy flaws, is **slick** in employing a "burden-shifting" tactic. This SEC's trick question is essentially saying, "We are already aware that the industry thinks the CAT is an illegal, unconstitutional mass-surveillance honeypot. Setting those arguments aside, what else do you have?" By asking if there are any other unaddressed concerns, the SEC builds a defensive record for future Court battles, where the SEC's legal counsel can argue "We gave the entire financial ecosystem a 60-day window to bring forward any hidden constitutional or privacy defects. If they didn't raise this specific technical vulnerability during the comment period, they cannot claim we failed to consider it."

The SEC exceeded its statutory authority by transforming CAT into a permanent, all-inclusive inventory of every American's financial actions. Reference to our response to Q9, despite the authorities' close oversight, the project migrated away from its explicit, post-2010 Flash Crash mandate – which was to create an anonymized, cross-market forensic tool – and instead evolved into a monolithic, un-vetted data repository. It is vulnerable to security threats, intruding upon privacy and **impairs the civil liberties** of Americans whom "transacting" or directly or indirectly

"engaging" in any way or form in U.S. securities markets.[20] The SEC and SROs have effectively built a **"CENSUS"** without the constitutional or statutory authority to do so.

(a) The Separation of Powers: Census vs. Securities Regulation

The U.S. Constitution explicitly delegates the power to conduct a national census to **Congress**, which has assigned that authority to the **Department of Commerce** (via the Census Bureau).

- **Strict Constitutional Boundaries:** The census has strict statutory rules, explicit limitations on what can be asked, and legal guarantees that individual responses cannot be used against a citizen by law enforcement or tax agencies.
- **The SEC's Usurpation:** Congress gave the SEC authority under the Exchange Act to regulate securities markets and prevent fraud – **NOT** to index the daily lawful behavior of millions of un-suspected citizens. By capturing every RFQ, order route, execution, and identity profile (via the CAIS), the SEC is conducting an ongoing, *unauthorized* **census** *under the guise of market oversight*.

(b) General Warrants and Mass Surveillance

The Fourth Amendment was written to outlaw "General Warrants" – instances where the Authorities searched entire towns or populations looking for a crime, rather than specifying a single suspect and a single location.

- **The Traditional Rule:** To see a citizen's bank account or trading records, the government historically had to show a judge "probable cause" or issue a targeted subpoena based on suspected wrongdoing.
- **The CAT Inversion:** The CAT flips this constitutional presumption on its head. It operates on a mechanism of permanent, automated **Massive Government Surveillance**.[20] It forces the collection of your data first, aggregates it into a central repository, and holds it indefinitely, allowing government observers to query your lawful activities at will.

(c) The SRO Monopoly as a Private Stasi

The laundering of government power through SROs and private corporations (FINRA CAT LLC/ Thesys Technologies LLC) is **unjust**. The SEC knew it lacked the direct statutory authority and financial resources to build this surveillance apparatus itself, allegedly in cahoots with, and commanded the SROs to build and pay for it.

- **Evading Constitutional Restraints:** By forcing SROs and the private CAT Plan Processor to collect PII (or the equivalent – CAIS) and order flows, the regulators attempted to create a loophole around the Bill of Rights, arguing that "private data collection" does not trigger Fourth Amendment protections.
- **SROs should NOT enjoy immunity related to their private business:** According to Weissman and Sparta Surgical Corp.'s Court cases against National Association of Securities Dealers (NASD),[80] FINRA and presumably all SROs remain subject to liability should claim(s) arise as a result of private business or commercial conduct. Given FINRA replaced Thesys Technologies (a private company) as the CAT Processor, indeed signified that FINRA and CAT LLC are in effect conducting private business. We argue such commercial conduct must be subject to corresponding risks and civil claims in the case of liability.[81]

(d) Shielding SROs from the Risks of an Accountable Dragnet

[80] https://caselaw.findlaw.com/us-9th-circuit/1274659.html
[81] https://www.linkedin.com/pulse/cat-outside-delegate-authorities-kelvin-to/

The SEC's failed attempt to slip industry-standard Limitation of Liability Provisions into the CAT Reporter Agreement (ultimately disapproved under Release No. 34-93484) lays bare the structural hypocrisy of the SRO monopoly.[37] Reference to our January 2021 comment letter,[82] we disagree with the SEC's commissioned Charles River Associates' Economic Analysis (CRAEA) – their three types of breach scenarios are insufficient to represent the potential damages to our country's economy and national security in case of breach.

- **The Sovereign Immunity and SRO Monopoly Contradiction:** The SROs operate as a commercial monopoly, yet they attempted to claim quasi-governmental immunity from cyber liabilities. If SROs and the CAT processor are granted an absolute shield against civil liability, it eliminates all market-driven incentives for them to maintain adequate data security. Forcing private market participants to hand over proprietary trading data under a mandatory regulatory dragnet – while simultaneously forcing them to sign away their rights to sue if that data is breached – is an **authoritarian overreach**.
- **CAT creates a systemic vulnerability that no commercial limitation of liability should ever protect:** The CAT operates under the naive assumption that a centralized data vault can be made perfectly secure through basic cloud controls. The SEC relied on CRAEA, naively capped estimated breach damages at a 95th-percentile loss of "greater than $100 million." This estimate grossly **undermines a potential multi-trillion-dollar National security threats**. CAT is a prime target for foreign adversaries. A breach is not a minor corporate loss; it could trigger a structural collapse of the U.S. capital markets. There are real-world precedents (see Q28 and Q29). Downplaying the risk of making a liability shield look acceptable was a dangerous misdirection.

The CAT is NOT a technical project that needs better cybersecurity; it is an unconstitutional overreach and deviated from the Congress's original mandate.

(e) Weaponization of the Audit Trail – Regulatory Overreach and Mission Creep

Under the guise of *"enhanced market oversight,"* the Commission and SROs are conveniently leveraging the CAT to bypass the formal legislative and administrative checks required for traditional rule changes (see Q1, Q16, Q20, and Q24). If the regulatory goal is to restructure market dynamics, tighten conditions on blocks, or alter institutional transparency, the Exchange Act demands distinct, independent rulemaking processes backed by explicit cost-benefit justifications. Instead, the SEC and SROs are weaponizing the CAT for commercial and political positioning, utilizing a mandatory data-vacuuming infrastructure to enforce sweeping policy changes without public accountability (see Q4, Q5, Q8, Q16, and Q25).

- **Subversion of the Legislative Mandate:** The CAT has radically departed from Congress's foundational focus on flash-crash prevention, morphing instead into an unauthorized, continuous "**census**" that operates entirely outside the SEC's statutory boundaries.
- **Circumvention of Independent Rulemaking:** Rather than pursuing transparent, individualized rulemaking to adjust conditions for e.g., Requests for Quotes (RFQs) or Large Trader 13(h) identifiers, the Commission uses the CAT as an administrative shortcut to expand its regulatory perimeter without proper procedural scrutiny.
- **The Section 31 Financial Transaction Tax:** Imposing the compounding infrastructure costs of this ever-expanding data vault onto the industry under the guise of *"user fees"* functions as an inequitable, backdoor contracts that suppresses capital formation and penalizes everyone like *Sh*t hit the fan*, see Q13 and Q15.

[82] https://www.databoiler.com/index_htm_files/DataBoiler%20SEC%20CAT%20Limitation%20Liability.pdf

- **Evisceration of Due Process via Blanket Surveillance:** While legacy procedures required specific thresholds or subpoenas to summon granular sub-account and beneficial shareholder data, the CAT demands the systematic, pre-emptive ingestion of proprietary lifecycle details without prior evidence of misconduct.
- **Commercial Exploitation by SRO Monopolies:** Private surveillance and software ventures tied directly to major Exchange Groups extract commercial utility from this aggregated market intelligence, forcing the rest of the industry to fund data pipelines that actively subsidize the SROs' private business segments. [1, 2]

Using Section 31 fee is tolling everyone in the industry, which will be passed-down to the end-investors. Should any information be obtained with higher urgency, or additional content is required, then the SEC must go through separate rulemaking. Conveniently pointing to the CAT to continually expand its scope to satisfy an unlimited desire to obtain anything the SEC and SROs want is unjust. Anything NOT following proper subpoenas or NOT tied to prosecuted cases and suspicious fraud or violation activities pertaining to "flash crash" or "market manipulations" ought to be removed from the CAT, see Q2.

27. **Should the Commission modify data-access controls** or other privacy protections? What changes and why?

The SEC / CAT "Secured Analytical Workspaces" (SAW) is problematic, see Q28. The CAT repository itself does NOT currently operate under a true Zero-Trust Architecture (ZTA).[83] While the SEC internally implemented a series of automated Zero-Trust safeguards in late 2024 to restrict how its own staff downlinks and shares CAT data externally, the structural repository managed by the SROs remains anchored to outdated, perimeter-based security controls.

(a) The Perimeter Myth vs. Zero-Trust Reality

The central CAT repository was built on a traditional *"castle-and-moat"* security philosophy.[84] It relies heavily on strict access controls, user provisioning, and the so-called *"Comprehensive Information Security Program"* (CISP) to guard the boundary.[85] See Q28 for our concerns with CISP.

The Flaw: True Zero-Trust operates on the principle of "never trust, always verify." It assumes the perimeter has already been breached. It micro-segments data, continuously authenticates every single transaction, and enforces end-to-end cryptographic borders.

The Current Gap: Because CAT relies on a legacy perimeter model, once a user or an SRO analyst is granted valid credentials to enter a *Secure Analytical Workspace* (SAW), they possess broad query privileges.[86] This design makes the database highly vulnerable to the automated, distributed query data leakage – like *Reconstruction and Model Extraction attacks*[87] – that Zero-Trust is built to disrupt.

(b) The SEC's Internal Asymmetry

The Inspector General and federal audits revealed that while the SEC rushed to deploy internal Zero-Trust controls to stop its own personnel from leaking data, the actual database architecture running at FINRA CAT did NOT follow suit. The regulator is attempting to secure its own terminal Endpoints while leaving the underlying central *"honeypot"* exposed to systemic infrastructure risks.[88]

[83] https://www3.weforum.org/docs/WEF_The_Zero_Trust_Model_in_Cybersecurity_2022.pdf
[84] https://www.cloudflare.com/learning/access-management/castle-and-moat-network-security/
[85] https://www.sec.gov/newsroom/press-releases/2020-189
[86] https://www.sec.gov/files/additional-oversight-monitoring-secs-cat-usage-needed-rpt-585.pdf
[87] https://arxiv.org/pdf/2502.16065v1
[88] https://www.cato.org/blog/should-consolidated-audit-trail-have-future

(c) Retrofitting Zero-Trust *"In the Air"*

The fact that the SEC is using the 2026 Concept Release to ask if they should mandate Zero-Trust proves that the system was launched without it. Forcing a massive, multi-terabyte live repository that ingests billions of market events daily to migrate to a Zero-Trust framework is highly complex. The SROs are already leveraging this difficulty to *push back*, pointing to the ballooning compliance costs and structural delays such a mandate would cause.

As I wrote in this 2023 article,[89] Zero-Trust is *"a shift in the security approach on how to dynamically and holistically establish trust with an unknown, whether a human being or a machine… Zero-Trust from a philosophical perspective could work in counter of people's social needs but is an important TECH bandage to curb bot attacks before a complete overhaul of the web… We agree there are significant gaps in the awareness and know-how in the appropriate implementation of Zero Trust Architecture. Yet, we believe 'Establishment Pushback' being the prominent factor hindering a paradigm shift."*

The Federal government's journey towards Zero-Trust was significantly accelerated by Executive Order 14028, which was issued in May 2021 and aimed to improve the nation's cybersecurity. It mandated Federal agencies to adopt ZTA to enhance their security posture and elevated the need for a unified approach that includes identity verification, device compliance, least-privilege access and continuous monitoring.

Federal agencies must implement multi-factor authentication (MFA) methods and grant access based on the principle of least privilege. This involves verifying the identity of users and devices before granting access to resources. Devices that access federal networks must be secure. This includes validating device compliance with security policies and continuously monitoring device health. Agencies must implement Endpoint Detection and Response (EDR) solutions to detect and mitigate threats at the device level.

Protecting sensitive data is a cornerstone of Zero-Trust. This involves encrypting data both *at-rest* and *in-motion* (see footnote 11) implementing Data Loss Prevention (DLP) solutions and continuously monitoring data access and usage. Zero-Trust requires the **segmentation of networks to limit the lateral movement of attackers**. *By dividing the network into smaller segments*, agencies can contain breaches and help prevent their spread across the entire network – addressing the *SAW broad privileges problem* mentioned earlier.

Reference to our cost estimate and response in Q23, we think patching legacy systems (OATS/ COATS/ EBS) is relatively cheaper and faster than the highly complex segmentation of the CAT system to migrate to a ZTA.

- $80-120 Million in near-term development Capital Expenditure, plus an ongoing annual operational overhead of $40+ Million. This funding is required to transition the central database into quantum-resistant encryption, maintain massive multi-tenant secure cloud enclaves, and administer continuous AI-driven penetration testing.
- **Caveats:** This security fix does NOT avert the risk of CAT being a prime target for sophisticated cyber warfare. As transaction volumes grow exponentially, the cost to store, link, and continuously scan a centralized repository scales linearly, creating a permanent funding bottleneck.

We understand the SEC has already been forced to spend millions and make dramatic changes because the CAT's security setup failed to meet modern standards. The legacy perimeter model that CAT relies on is highly vulnerable to the automated, distributed query data leakage. The CAT repository processes billions of algorithmic data messages every day. Securing a database of that size against modern cyber threats requires continuous, expensive upgrades.

[89] https://www.linkedin.com/pulse/improving-trust-amid-race-technologies-kelvin-to-8vxrc/

Funding constant patches for cloud-based secure enclaves, maintaining strict field-level masking, and running continuous automated penetration tests creates a perpetual financial drain – i.e., a **money pit**.

It is unrealistic to fulfill unlimited desires – wanting CAT to be "everything everywhere all at once" to everyone. **The CAT is broken beyond repair** and is wholly inadequate for the challenges of the twenty-first century (see Q3). The ONLY way to reduce the CAT's footprint is right coursing its scope and architecture. Instead of choosing between keeping OATS/ COATS/ EBS or CAT, the Commission should make a bold move to *shift away from continuous trade reporting* (in CAT or these legacy systems). Instead of forcing market participants to constantly transmit billions of raw, duplicate logs to a central *honeypot*, the data would remain locked at its native origin: inside an exchange's matching engine, a clearing house's ledger (like the DTCC or NSCC), or an SRO's internal database. Regulators would instead deploy automated*, federated analysis models* directly to those environments (see Annex 2).

Section H — Cybersecurity

28. **Should the Commission modify cybersecurity requirements** for CAT, EBS, LOPR, or other audit trails? What changes and why?

This SEC concept release reopens fundamental questions regarding cybersecurity, data privacy, governance, and structural scope of the CAT, EBS, and Large Options Positions Reports (LOPR). Asking this question at this time is like the plane already taking flight while the engineers are still trying to figure out how to assemble the parts in the air. The trading and investment communities have repeatedly warned that User-Defined Direct Queries (UDDQ) and bulk extractions allow sophisticated actors to reverse-engineer proprietary algorithmic models based on data in the repository. Attempting to bolt structural cyber defense onto a live repository processing billions of daily market events is fundamentally more volatile than embedding privacy-by-design from inception.[90]

As proposed by the SEC, the CAT Operating Committee established the Security Working Group (SWG) for CAT. SWG membership roster excludes frontline industry market participants, it functions primarily to protect the *repository's perimeter* using ONLY standard infrastructure controls. Per our November 2020 comment letter,[91] we already informed the SEC that their proposal[92] referenced an outdated revision 4 of NIST's SP800-53. Sadly, an April 28, 2026, FINRA CAT regulatory filing reveals that their CISP continues to rely on this over a decade old revision 4.[93]

Inadequate Protection for Cloud Ingestion: Revision 4 was not built to counter modern cloud-native exfiltration threats associated with unconstrained, multi-tenant cloud-buffering loops. This gap leaves massive real-time data aggregates vulnerable to insider threats and sophisticated supply chain intrusions.[94]

The "Bare Minimum" Compliance Lag: This rigid alignment incentivizes organizations to clear the minimum bar of an outdated standard rather than proactively adopting modern best practices. The SEC and SROs escape accountability for their own technical stagnation while aggressively burdening the industry with stringent cybersecurity disclosures. Strikingly, the Commission itself has been the subject of multiple high-profile security failures.[95]

The Policy "In-Between Time": The time required to draft, notice, comment, and implement regulatory policy changes creates an operational vacuum. Emerging AI-driven extraction methods evolve in days, while formal standard updates take years, leaving data exposed to novel threats during the regulatory lag.

The standard for secure cloud environments has evolved significantly. The Federal Financial Institutions Examination Council (FFIEC) formally sunset its own old Cybersecurity Assessment Tool framework. Banking and credit regulators actively push financial institutions toward NIST CSF 2.0 because outdated, descriptive checklist models are no longer sufficient to secure modern data systems. By testing the CAT data vault against older security framework parameters, the SWG overlooks the critical components embedded in newer standards:

The Absence of Advanced "Govern" Protocols: Modern frameworks like NIST CSF 2.0 introduce dedicated **Govern** functions that demand strict, transparent, board-level oversight and aggressive supply chain risk mitigation. The current SRO-dominated governance model avoids this multi-stakeholder accountability.[96]

---

[90] https://www.sifma.org/issues/regulatory-compliance/consolidated-audit-trail
[91] https://www.databoiler.com/index_htm_files/DataBoiler%20SEC%20CAT%20Enhanced%20Security.pdf
[92] https://www.sec.gov/rules/proposed/2020/34-89632.pdf
[93] https://www.finra.org/sites/default/files/2026-04/SR-FINRA-2026-011.pdf ;
[94] https://www.govinfo.gov/content/pkg/FR-2026-05-12/pdf/2026-09342.pdf
[95] https://cybernews.com/cybercrime/hacker-guilty-sec-account-hack/ ; https://www.sec.gov/newsroom/press-releases/2017-170
[96] https://www.myabt.com/blog/ffiec-cat-nist-csf-2-financial-institutions

Concerns highlighted in our November 2020 have NOT been addressed and they are still valid, let us revisit.

> It is interesting that while the Commission allows themselves to attend all meetings of the SWG as observers, when there is a likelihood that SWG may review system log records to assess if there might be potential abusive usage of CAT by allegedly Commission staffer. But then, the Commission limits the SWG access to CAT data per footnote 33 of the SEC's proposal (Release #: 34-89632).
>
> To earn public trust on CAT, the first priority is to allow the public to scrutinize government agencies and SROs' actions, not the other way around. If SWG and CAT Operating Committee are representative of public interests rather than just "*participants*" from the Exchanges, they should have rights to challenge anything they deem *inflicting damage on* "*others*" as a result of *inappropriate use* or *negligence in the design* of CAT.
>
> We are not sure what constitute "implements all common technical security controls required by the CISP" and how it would be enforced. The nine groups of participants may each have their own SAW. Also, security and privacy controls would always be a race against hackers, by setting a minimum standard rather than pursuing the best defense would introduce opportunities for the hackers. By the time the participants agreed to a *"common"* SAW (see Q27), the techniques may have been made obsolete.
>
> Once a hacker is inside SAW, it is up to the other controls to make it sufficiently hard or to not enough time for the hacker to inflict any damage, as well as pessimistically, to minimize loss when all else fails. Thus, even inside SAW, we recommend protocols to address security risks for *data-in-motion*, *in-use*, *at-rest*, *disposal*, and restrict the volume and frequency of data queries for *"specific case investigation"* or *"rule enforcement"* purpose, etc.

It is not just about expertise in information security and privacy, but also ethical/ civic concerns that the SEC and SROs ought to address. Continuous building of data controls, and governance policies and procedures may unintentionally create big bureaucracy overtime. The minute the CISP or audit trail cybersecurity requirements are formalized, it is the minute its enemies – foreign adversary / bad actors would know and would attempt to hack or break it. Without embedding appropriate analytical frameworks into the design of the CAT, as we have pointed out since our comments in 2016, CAT is a useless, gigantic vault that does nothing other than cause disturbances to all industry members wasting valuable time and energy in data submission and causing worry about security and compliance.

The CAT's CISP lacks robust scenario planning or defensive engineering – such as differential privacy or cryptographic boundary defenses – to block distributed learning attacks designed to reconstruct underlying order flows. The CAT is completely devoid of AI defenses. Our previous identified technical vulnerabilities – Membership Inference, Reconstruction, Property Inference, and Model Extraction Attacks – remain unaddressed.

Unlike the **NSA, DHS, or the National Cyber Director**, the SEC is a disclosure and enforcement regulator, not a specialized defense agency. The Commission has ignored our past repeated reminders on security. The Edward Snowden Case proves that even highly classified intelligence systems are vulnerable to trusted insiders and mass data exfiltration.[97] The SWIFT Banking Hack demonstrates how hackers can weaponize stolen credential data from a trusted entity to execute multi-million-dollar thefts against central banks.[98] The SEC undermined these threats.

From the get-go, the Commission should NOT attempt to prescribe any particular form of security control measures or best practices. Instead, there should be high-level principles to guide the CAT project to embed essential security, privacy, and analytical frameworks in its design.

[97] https://en.wikipedia.org/wiki/Edward_Snowden
[98] https://en.wikipedia.org/wiki/2015%E2%80%932016_SWIFT_banking_hack

Please revisit our "**A through Z**" security and privacy suggestions.[99] We reiterate our response to Question 11 of our November 2020 comment letter, *"As long as CAT remains a gigantic vault, access to the vault would always be a complicated question. One can set up the best structural controls, multi-layers of defenses and what not. Yet, hackers are equally familiar with these structural parameters, if not smarter. A well-organized and standardized approach to privacy and security control is NOT better than data obfuscation through random. So, an alternate approach to security control is called 'chaos' (i.e. without a known structure). As an analogy, port cities would make their streets like a maze to deter pirates and other adversaries. The more organized things are, the easier it is for hackers to crack down patterns in breaching controls. Chaos or 'unknown unknowns'* would significantly heighten the difficulty of access. Remember, hackers steal what is easy and *'common'*."

29. **Should the Commission require additional cybersecurity testing, audits, or certifications?** What requirements and why?

It is a tech race, the so called *"testing, audits, or certifications"* only benefit the bureaucrat, like those lawyers' billable hours on CAT. It is about ensuring appropriate cyber defenses are embedded in the design. A design flaw can NEVER be addressed by more testing, audits, or certifications.

[99] https://www.linkedin.com/pulse/cat-through-z-security-privacy-requirements-kelvin-to/

Section I — Transparency and Process of Comprehensive Review

30. **Should the Commission adopt additional transparency measures** regarding CAT operations, costs, or performance? What measures and why?

The SEC has maintained a formal representative observer seat on the CAT Advisory Committee and has been embedded in the operational process since Rule 613 was first finalized. Despite this close oversight, the project migrated away from its explicit, post-2010 Flash Crash mandate—which was to create an anonymized, cross-market forensic tool—and instead evolved into a **monolithic, un-vetted data repository** processing the PII (or the equivalent – CAIS) and daily trading intellectual property of millions of Americans, see Q9.

When a database is fundamentally flawed by architecture, adding "transparency measures" or altering reporting formats will not resolve its underlying structural and constitutional vulnerabilities. For example, the trading and investment communities have repeatedly warned that UDDQ and bulk extractions allow sophisticated actors to reverse-engineer proprietary algorithmic models and other issues. In the last decade, we at Data Boiler informed the SEC about the wrong reference to an outdated version of NIST SP800-53 and suggested various improvement across various aspects of CAT, but we found ourselves **talking to a wall, without any substantive fixes or concrete actions taken thus far**, except disapproving CAT limited liability provisions for SROs.

Transparency initiatives are and will continue to fail to change the core trajectory of the CAT. Any transparency rules the SEC passes merely command the SROs to publish more detailed performance metrics, audit results, or budget breakdowns. Making an inherently insecure or overly broad collection practice more transparent does not change what data is being ingested. Increased disclosure regarding how data flows inside the system does NOT fix the **Fourth and Fifth Amendment overreach** that civil liberties groups and industry participants continue to challenge in the Federal Courts, see Q26.

Transparency simply gives the public a clearer view of the apparatus that looks well decorated outside without offering any structural mechanism inside to halt it. By acting as the primary driver behind the CAT repository while lacking the institutional cybersecurity capabilities of the NSA or DHS, the SEC created a systemic single point of failure, see Q28.

For years, the industry was fully aware of the CAT's exploding budget. Little technology development works were capitalized. CAT's annual operating budget ballooned from an initial projection of **$55.8 million** to somewhere between **$188 million** and **$272 million**, see Q13. Recent regulatory interventions forced the SROs to find approximately 100 million in projected budget rollbacks, but this amount will have little effect in averting the projected hyper growth in storage costs. The sheer technical design of the CAT has morphed into a **bottomless data warehouse** that **ballooned** from a historic average of 296 billion to the recent average of 700 to 800+ billion records every single day.

Knowing the exact price tag of the repository does NOT fix its operational vulnerabilities and would NOT avert the harsh reality that the **CAT is unsustainable**. Following the Eleventh Circuit Court of Appeals vacating the 2023 Funding Order, subsequent lawsuits by Citadel Securities to freeze reserves, and a projected "coffer depletion date", the CAT faces an existential structural crisis. The CAT centralized database remains a consolidated *honeypot*, regardless of what the SEC and SROs are, or will, do to it unless they redefine its scope and architecture.

Publishing charts on data ingestion errors or system latencies does NOTHING to embed **privacy-by-design** into a system that was built from the ground up without it, see Q27. It is unrealistic to fulfill unlimited desires – wanting CAT

to be "everything everywhere all at once" to everyone, see Q16. Focusing on performance metrics completely overlooks the primary structural vulnerability: a centralized data vault that fundamentally compromises ZTA.

31. **Should the Commission modify the process for conducting comprehensive reviews** of audit trails and data sources? What changes and why?

How this Apr 2026 Concept Release for a "comprehensive review" of the CAT insufficient in documenting the evidence in supporting the Commission on whether to keep or kill CAT/ OATS/ COATS/ EBS in the short- and long-term once-and-for-all? This question indicates that the SEC may want to run additional reviews, or make this exercise an ongoing routine, where if it is true, it would be burdensome. Kicking the can down the road is the worst move compared to a decisive stance on whether to keep it or killing it.

This flawed "elephant" was given over a decade to experiment, not once (Thesys) but twice (FINRA CAT LLC), to do a $2+ billion proof-of-concept about a "centralized single source of truth" that is doomed to failure. The industry cannot afford *frequent pestering* from the SEC and SROs wanting to make modifications here and there, patch work changes, and run ever more experiments with CAT over and over again. The SEC/ SROs' unlimited desires – wanting CAT to be "everything everywhere all at once" to everyone, is harming everyone. Conducting *"further reviews"* is an institutional stall tactic. Procrastination allows the status quo to crystallize, benefiting the cloud vendor. The longer the CAT operates in its current vulnerable state, the more *"entrenched"* it becomes, making it harder to pull the plug.

The SEC is structurally blindsided. A massive asymmetry of technical expertise exists that hopefully the Commission leadership would have been educated on by now. Immediately, the SEC, GAO, and DOJ must investigate if CAT data was **weaponized for political or commercial reasons** rather than serving the Congress's mandate. Or else, by the time the review concludes in a few years, billions more market events will have been ingested, and the SROs will argue that the system is too integral to market structure to be structurally altered.

Public confidence in market integrity erodes as the SEC continues to tolerate chronic national market disorders and data distortions within CAT and SIP. This instability creates a chilling effect in discouraging market participation and places undue burden to market participants. In turn, capital will migrate toward digital assets, gaming platforms, and offshore jurisdictions, jeopardizing the dominant position of the U.S. financial system.

Even the SEC and the best talents attending CAT meetings and *"providing views"* means nothing. The CAT Operating Committee can politely listen, check the regulatory compliance box, and then vote to execute whatever architecture serves their corporate interests. The financial architecture of the CAT represents a textbook case of moral hazard. Because the SROs designed a funding model that passes the vast majority of CAT costs directly down to broker-dealers (and ultimately, retail investors), they face zero economic pressure to build something efficiently or securely.

The project migrated away from its explicit, post-2010 Flash Crash mandate—which was to create an anonymized, cross-market forensic tool—and instead evolved into a **monolithic, un-vetted data repository** processing the PII (or the equivalent – CAIS) and daily trading intellectual property of millions of Americans, see Q9. The question is NOT about who is who on which committee/ groups; who was, or is, the CAT Plan Processor that will or will not listen to advice, but the **faulted and outdated design of CAT**.

Again, we would like to reiterate – CAT LLC could have directly aggregated these high-speed, native pipelines to augment the missing order-level details that legacy systems like OATS and MIDAS lacked. Had regulators simply synthesized these existing proprietary feeds alongside centralized clearing and settlement systems, the SEC and SROs could have conducted comprehensive market-wide analysis without ever needing to construct a centralized, multi-billion-dollar CAT data vault (see our response to Q4 and Annex 2).

Section J — General Request for Comment

32. **Are there any other considerations, costs, burdens, or benefits** related to audit trails or data sources that the Commission should evaluate?

The flaws of this outdated design of CAT are widely known (e.g. the 50± millisecond timestamp tolerance, the missing of futures and swap data, cybersecurity threats, civic concerns for massive government surveillance, etc.), but few dare to say because of worry about *retaliation*. We beg all conscience authorities to stop this money pit immediately and revisit the CAT design.

According to SIFMA, the total private-sector compliance expenditure directly incurred by industry members to report data has exceeded $1.7 billion annually. Industry members have already paid more than their fair share for CAT. We do NOT believe control measures, such as monitoring and metering the ad-hoc queries of enormous data sets would move the needle in shaving "fat" off this "outsized elephant", see Q12. Cost-benefits analysis should NOT ignore this significant social cost incurred by industry members, see Q13, Q15, and Q26. Rulemaking to seek sole benefit for the government agency or the affiliated SROs should be prohibited because it is in contrast with serving the public interest.

The sheer technical design of the CAT has morphed into a *bottomless data warehouse that balloons* from an historic average of 296 billion to recent average of 700 to 800+ billion records every single day. This flawed "elephant" was given over a decade to experiment, not once (Thesys) but twice (FINRA CAT LLC), to do a $2+ billion proof-of-concept about a "centralized single source of truth" that is doomed to failure. Following the Eleventh Circuit Court of Appeals vacating the 2023 Funding Order, subsequent lawsuits by Citadel Securities to freeze reserves, and a projected "coffer depletion date", the CAT faces an existential structural crisis.

Reference to our cost estimate and response in Q23 and Q27 we think patching legacy systems (OATS/ COATS/ EBS) is relatively cheaper and faster than the highly complex segmentation of the CAT system to migrate to a ZTA. **Caveats:** This security fix does NOT avert the risk of CAT being a prime target for sophisticated cyber warfare. As transaction volumes grow exponentially, the cost to store, link, and continuously scan a centralized repository scales linearly, creating a permanent funding bottleneck. Instead of choosing between keeping OATS/ COATS/ EBS or CAT, the Commission should make a bold move to *shift away from continuous trade reporting* (in CAT or these legacy systems).

It is unrealistic to fulfill unlimited desires – wanting CAT to be "everything everywhere all at once" to everyone. **The CAT is broken beyond repair** and is wholly inadequate for the challenges of the twenty-first century (see Q3). The ONLY way out is right coursing the scope and architecture of CAT (or a new forensic tool). See Annex 2 to marvel at the advancement of Agentic AI for better ways to detect, hunt down bad actors/ foreign adversaries and prevent flash crashes while respecting civil liberties, improving privacy and securities, at a lower cost and reduced footprint.

33. **Are there potential regulatory responses not identified** in the release that the Commission should consider?

This question implicitly means – "Following the Eleventh Circuit Court of Appeals vacating the 2023 Funding Order, what other legal or structural mechanisms can we deploy that would not result in another immediate lawsuit?" We observe that the SEC seems to be running out of viable ideas to salvage the CAT's current trajectory. We strive to provide objective analysis that illuminates operational realities, system vulnerabilities, architectural alternatives, and governance issues—addressing critical areas where the Commission may otherwise lack visibility. Please excuse our honesty in bluntly citing areas of concern.

We recommend shutting the CAT off immediately, patching OATS/COATS/EBS in the short-term, and go back to the drawing board – Synthesized existing proprietary feeds alongside centralized clearing and settlement systems. Bad

actors / foreign adversaries play across markets and payment systems simultaneously. All five Dodd-Frank regulatory agencies (SEC, CFTC, OCC, FRB, FDIC) must break down silos and work in unison to safeguard US financial stability. Modernized Cross-Market and Payment Platforms monitoring with Agentic AI to assess changes in market dynamics (convergence of TraFi and DeFi), to understand where frictions and liquidity concentration may shift at rapid pace or irregularities and dysfunctional market behaviors occur. Analyze the stress points and weak links where the next **flash crash** or systemic risks may emerge, and develop mitigation protocols with more effective and timely volatility interruption mechanisms.

While many understandably worried about AI displacement, they must nevertheless learn to orchestrate these new tools. *A good decision, made now and pursued aggressively, is superior than a perfect decision made too late.* Humans are slow and CANNOT manually reconcile the massive volume of structured trade logs and unstructured data driving modern markets. AI bridges this gap by handling tedious data ingestion and synthesis. Supported by human context, institutional knowledge, and strict guardrails, AI acts as a force multiplier – not job replacement. This shift elevates agency personnel from manual data processors to strategic gatekeepers of market integrity.

See Annex 2 in this comment letter for an elaborate discussion of our innovative design to make the CAT replacement nimble and effective with agentic AI, Model Context Protocol (MCP) to connect and orchestrate workflows across systems, Retrieval-Augmented Generation (RAG) for knowledge retrieval to avoid wasting tokens on hallucinations, and surround all with Zero-Trust cyber defense to safely manage what the AI can see and do. Collectively these components form the "layers of intelligence" to address the challenges of today and the future.

Finally, the SEC, GAO, and DOJ should investigate whether the CAT data was **weaponized for political or commercial reasons** that benefit the SROs and their vendors.

Disclaimer: Nothing contained in these submitted comments shall be construed as providing legal advice or establishing an attorney-client relationship, as we are not attorneys and do not offer legal counsel in any way, shape, or form. These comments are submitted in good faith as a matter of public interest, exercising protected rights of speech and regulatory petition; under no circumstances shall this submission expose the author to civil liability, defamation claims, or retaliatory legal action by any party who may disagree with or object to the critiques expressed herein. Furthermore, under no circumstances shall this submission be interpreted as permitting the SEC, the SROs, or any other party to shift the operational, architectural, or intellectual burden of securing and organizing this multi-billion-dollar database away from the Commission and onto the public taxpayers and market participants who are forced to fund the CAT project.

ANNEX 2 – Data Boiler's Vision of Next-Gen Market Monitoring System to replace CAT

CAT is a monolithic, un-vetted data repository. It is vulnerable to security threats, intruding upon privacy and impairing civil liberties of Americans who are "transacting", or directly or indirectly "engaging" in any way, shape, or form in U.S. securities markets.[20] Framing CAT's replacement as a choice between developing a different audit trail or relying on legacy data sources is a false dichotomy; far more innovative alternatives exist. Regulators and SROs do NOT need another centralized audit trail to do their jobs. Digital trails already exist publicly, albeit largely in unstructured forms. The focus must shift away from collecting data as the answer; what is truly missing is the intelligent analytical layer.

Ever since the CAT was introduced over a decade ago, market participants have stressed the need to analyze the interplay between securities and CFTC-regulated Designated Contract Markets (DCMs) and Swap Execution Facilities (SEFs). Also, multi-asset convergence between TradFi and DeFi is inevitable. New risks have emerged. Traditional guardrails may merely displace fraud and manipulation risk from SEC-regulated markets, rather than preventing that risk from proliferating elsewhere. Bad actors / foreign adversaries play across markets and payment systems simultaneously. Regulatory market monitoring ought to be modernized alongside market and technological advancements.

To create an anonymized, cross-market, cross-asset class forensic tool in fulfilling the post-2010 Flash Crash Congress's mandate, CAT's replacement must be nimble and effective with **agentic AI**, Model Context Protocol (**MCP**) to connect and orchestrate workflows across systems, Retrieval-Augmented Generation (**RAG**) for knowledge retrieval to avoid wasting tokens on hallucinations, and surround both with **Zero-Trust** cyber defense to safely manage what the AI can see and do. Please see below for our schematic drawing, collectively these components form the "layers of intelligence" to address the challenges of today and the future.

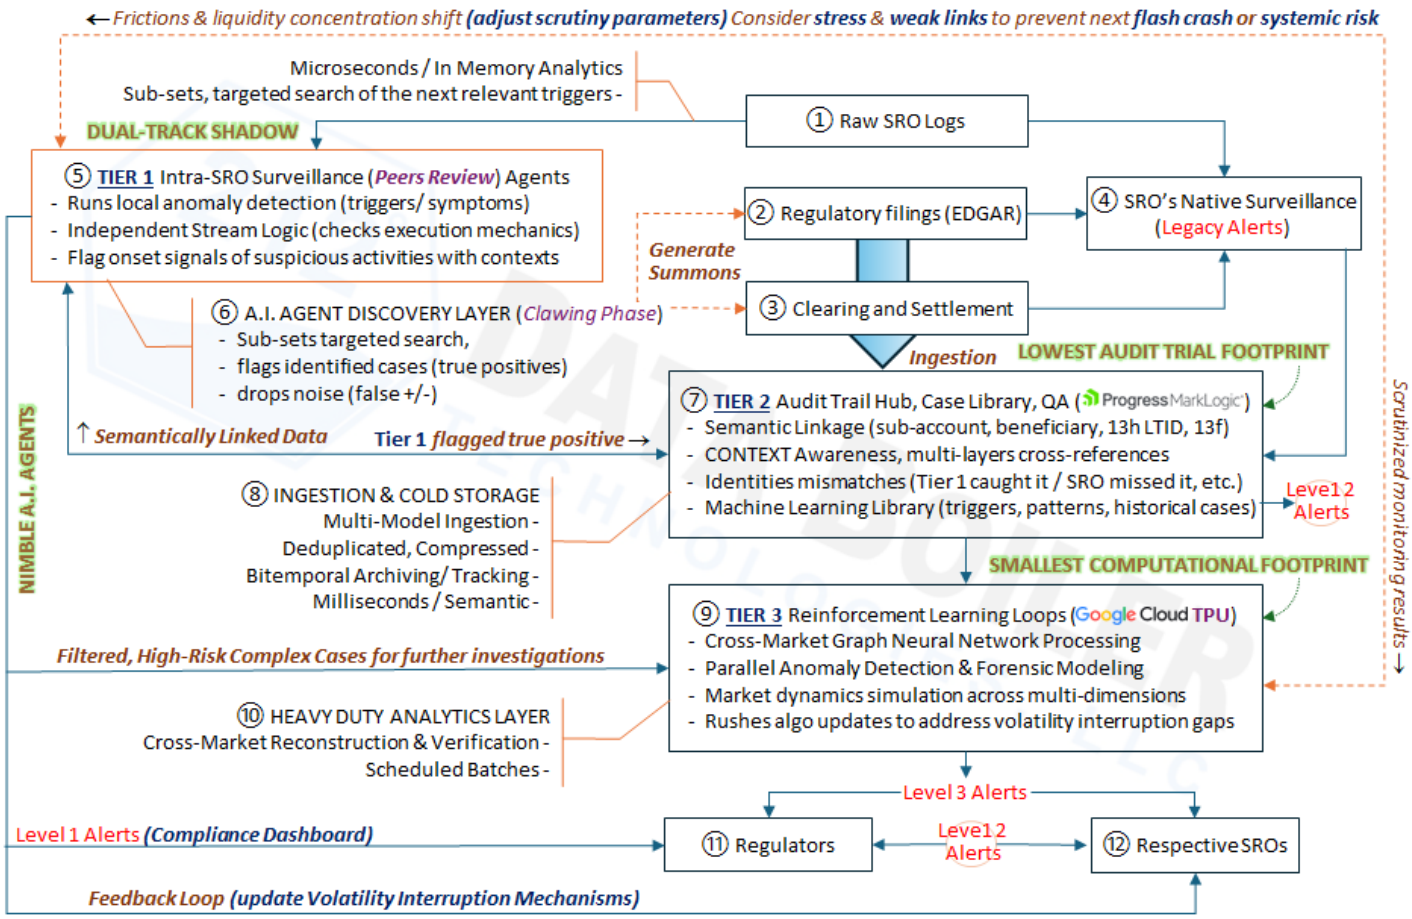

Goals

(a) **Lowest Audit Trail Footprint:** Drastically reduce Cloud Hosting Services cost, no more CAIS, minimize data-in-motion, only suspicious activities and confirmed and past prosecuted cases are stored in the Zero-Trust secured library.

(b) **Smallest Computational Footprint:** No finding a needle in a haystack, context awareness to pull ONLY the relevant, orchestrated AI agents for concerted efforts in fabricating layers of intelligence, minimize rework/ recalculation, and be scalable with the best performance-per-dollar.

(c) **Strategic gatekeepers of Market Integrity instead of manual data processor:** No more tedious treatment (insertion/ deletion of nodes, data labeling, corrections) of data, no more CPU-bound data-wrapping process to regurgitate SIP and other public data, avoid noise introduction, no more cumbersome representative order linkage, no more multiple sequential API calls to aggregate data from different sources that bloat the context window, no more writing tedious translation layers to accommodate quirks, different formats, and strict schema.

(d) **Get rid of unnecessary trade reporting:** relief of broker-dealers' burden, streamline and automate the procedure to summon sub-account and underlying beneficial owner information, eliminate and bypass downstream data intermediaries to analyze information directly at the source, deduplicate overlapping data streams, and selectively ingest only high-value, enriched contents.

(e) Optimize performance (more accurate detection of onset signals / anomaly at accelerated speed to generates alerts and reduce false ±) and be effective to assess changes in market dynamics, understand where frictions and liquidity concentration may shift at rapid pace or irregularities and dysfunctional market behaviors occur. Analyze the stress points and weak links where the next **flash crash** or systemic risks may emerge, and develop mitigation protocols with more effective and timely volatility interruption mechanisms.

The Meta-Regulatory Agent Topology

Unlike the monolithic, un-vetted CAT data repository that is vulnerable to security threats, intruding upon privacy and impairing civil liberties of Americans who are "transacting" or directly or indirectly "engaging" in any way, shape, or form in U.S. securities markets, our design is a **Distributed/ Federated, Event-Driven Agent Topology**. Given over 150 trillion market events annually, breaking up the surveillance routine into – **TIER 1** Intra-SRO Surveillance (***Peer Review***) Agents → ⑤ in the drawing; **TIER 2** Audit Trail Hub, Case Library, Quality Assurance for macro-filing lookup → ⑦; and **TIER 3** Reinforcement Learning Loops for cross-market Tensor Processing Unit (TPU) verification → ⑨ to achieve the above stated goals.

TIER 1 *Intra-SRO Surveillance (**Peer Review**) Agents*

Instead of forcing a single, vulnerable central vault to act as an omniscient *"brain"* trying to look for "everything everywhere all at once", this Tier 1 approach uses a "**Dual-Track Shadow Processing**" layout to independently verify an Exchange's surveillance system effectiveness at the edge. As raw matching engine logs stream in, they split into two paths: Track A goes through ④ the SRO's legacy surveillance tool that consumes ① raw logs from SRO's matching engine + ② regulatory filings (e.g. 13F/ 13H filings in EDGAR system) + ③ data in clearing and settlement systems as their inputs, while Track B simultaneously connects securely into an independent, localized Tier 1 AI agent for real-time analytics.

Much like an HFT firm algorithmically isolating *toxic versus natural liquidity*,[100] or an ATS operates as a midpoint-only venue using continuous, AI-driven price optimization to minimize *adverse selection*,[101] this layout capitalizes on existing,

[100] https://patents.google.com/patent/US7587347B2/en ; https://patents.google.com/patent/US7987128B2/
[101] https://patents.google.com/patent/US20210272201A1/en

specialized technology. Rather than forcing the industry to reinvent the wheel for a centralized government database, the knack here is leveraging "*multiple distributed brains*" to conduct continuous, **automated peer review** right at the source.

The Commission's recent proposal to rescind Rules 611 and 610(e) of Regulation NMS[102] provides an alternative path to architectural design. For two decades, virtually all trading platforms and SROs have maintained highly complex, ultra-low-latency technical stacks – consuming a synchronized mix of proprietary feeds and the SIP – to comply with federal trade-through bans[6] and locked/crossed market restrictions.

As the Commission moves to dismantle these legacy mandates, these massive industry tech investments do not need to become stranded, obsolete capital. Instead of forcing the industry to fund a redundant, centralized CAT *honeypot*, these highly sophisticated, pre-existing routing and data-ingestion architectures can be seamlessly realigned and redeployed. They are perfectly positioned to serve as the local infrastructure foundation for the independent, edge-based analytics required by the Tier 1 Agentic system.

To make the ⑤ Tier 1 detection engine even more nimble and agile, we can break a full scheme of market manipulation "pattern" into digestible pieces. Typically, deceptive market manipulation typically comprises characteristics resembling a mix of, but not limited to, the following *"triggers"*:

(a) continuous alpha trade with large accumulated losses;

(b) unusual trade volume (rigged-up volume with precipitous price fall);

(c) inexplicable price run-ups, run-downs, or spike(s);

(d) suspicious mark price near time of market close;

(e) trade style inconsistency with quick reversal of direction;

(f) frequent buy and/or sell among the same business group or affiliates;

(g) dealing with a concentrated group of dealers;

(h) suspicious price speculation on or around an initial public offering (IPO) date;

(i) infrequent trade, small-cap asset irregular activities (for small-cap securities where volumes are typically not high, a sudden large number of shares being traded may be a sign of market manipulation);

(j) cumulated outsized positions with size that is greater than a configured percentage of supply;

(k) dominated long position with series of heavy sell orders;

(l) off-load substantial poor performing asset to retail;

(m) retail rarely squared-off their pumped positions;

(n) erosion of market liquidity with raising of perceived liquidity;

(o) suspicious off-street financing corner with a squeeze;

(p) flooding of market with devalued collaterals;

(q) wave of margin calls in case of a sell-off;

(r) suspicious set price using long call with synthetic call;

(s) trading glitch with algorithm deployment error or control being off;

(t) risk, volatility and/or market prices moved with destabilization;

(u) tight two-way price quotes with a wave of cancelled trades;

(v) substantial orders placed with quote amendments.

[102] https://www.sec.gov/files/rules/proposed/2026/34-105655.pdf

Any single "trigger" listed above may not be a violation. However, the legitimacy of trade activities may be challenged if they occur in an order resembling a set of triggers (i.e. a lesson) in the machine learning library, see below table:

#	Manipulation	Pattern	Flag / skew scores
(i)	Corner/ squeeze	jbeno	Stop trades that will impair clients with squeezed prices
(ii)	Pump and dump scheme	bcine	Stop heavy sell orders
(iii)	Boiler room tactics	lim bcne	Stop further sales of poor performing assets to retail
(iv)	Paint the tape/ wash trade	fbd	Stop further sales of poor performing assets to retail
(v)	Circular trading/ primary dealer scheme	ghica	Stop trades that will impair clients with discriminated prices
(vi)	Forced liquidity scheme	kjunot	Stop strong buy orders
(vii)	Stock bashing/ reverse corner/squeeze	kpcq	Stop strong buy orders
(viii)	Manipulate price to prevent margin call	arce	Stop the related call option / synthetic call order
(ix)	Spoofing/ layering/ stuff quote	vuce	Stop further cancel/amend activities
(x)	Marking the close	fgdc	Stop trades that will destabilize price
(xi)	Manipulated trade spikes	scuv	Stop quick reverse direction trades

As illustrated above, the segmentation and grouping of trade styles or patterns helps to expedite the comparison process in the comparative analysis, by avoiding the need to compare against all the lessons in the Machine Learning Library. The segmentation and grouping of trade styles or patterns can also represent market change signals, suitable for Flash Crash prevention. Historical cases can be used as Complex Event Processing (CEP) lessons stored in the ⑦ Machine Learning Library (**Progress MarkLogic** semantic database in Tier 2) to train the Tier 3 TPU Reinforcement Learning model, to improve the ability of onset signal detection to recognize the symptoms of market manipulation, market change signal, synthetically created trades, or likelihood that the set of trade activities will result in a market price move of one or more financial assets against plans, or other signals.

A matching process is performed, which compares the ① Raw SRO logs like a stream of music tracks with the *"triggers"* stored in the machine learning library (Tier 2 semantic database). When the pattern of these logs/ tracks resembles the pattern of a trigger, a trigger counter is updated. The trigger counter is used to identify a subset of the lessons which may match a complete lesson, where subsequent comparisons are performed ONLY for the subset.

For example, when the trigger "j" is matched, the *trigger counter* for "j" is updated. The matching process is able to determine from the updated *trigger counter* for "j" that the lessons (i) and (vi) are possible matches while the others are not. The matching process then proceeds to only compare the subset of lessons ⑥ and foregoing comparisons with the remaining lessons, thus optimizing the matching process. The score for the trade activity is adjusted with increasing weight with each match to a trigger in a lesson. For example, the score is adjusted in response to a match to trigger "b" in lesson (i). As the matching process subsequently matches trigger "e", "n", and "o", the score is adjusted in increasing weights, i.e., reflect increasing likelihood of trade irregularities, where the match of "o" has the most weight.

This saves tremendous computing power and is more effective than the current process of finding needle in a haystack with manual queries. Also, these *"triggers"* help to form the basis upon which to determine if a combination of trades may become

synthetically created trades, particularly useful for subsequent Tier 3 TPU Reinforcement Learning for Cross-Market and Cross-Asset Serveillance.

The Tier 1 local agent behaves as an MCP Server uses Local Standard Input/Output (Stdio) Transport for ultra-low latency and maximum security, ensuring data never travels over a public network prematurely. When a query hits Tier 1, the MCP Server maps the request against the later discussed Tier 2 **Progress MarkLogic**'s Element-Level Security (ELS).[103] The Tier 1 agent generates a standardized data package containing ONLY findings alongside the metadata of the Level 1 Alerts generated (or ignored) by the SRO's native system. These local agents operate incredibly fast using minimal, temporary memory cache. Hence, the payload output is optimized.

TIER 2 *Audit Trail Hub, Case Library, Quality Assurance (**Semantic Context Awareness**)*

If an order and trade sequence stream behaves completely normally as determined by ⑤ the Tier 1 Agents, the data bypasses heavy evaluation and is compressed in a relatively cheaper storage ⑧. When an anomaly is triggered, semantic database such a **Progress MarkLogic** in ⑦ Tier 2, would unpack that data and index it as an interconnected network of relationships making semantic inferencing (figuring out hidden relationships between accounts, entities and trades) readily available for subsequent further investigations. While the components in ⑦ handles the database schema and stores the semantic links, it does not possess the native mathematical power to run complex machine learning heuristics over billions of data point, that is where the ⑨ Tier 3 TPUs come in.

The fundamental challenge of a CAT redesign is data harmonization. SRO matching engines output dense, fast-moving, and rigid streaming events, whereas regulatory filings are unstructured or semi-structured documents. **Progress MarkLogic** is a **multi-model NoSQL document database** featuring native XML/JSON document storage, an integrated triple store for semantic knowledge graphs, full-text search, and robust bitemporal capabilities. **Progress Semaphore** adds the ontology management and semantic enrichment layer on top, described below. This unique toolset addresses several key operational requirements:

- Handling Schema Drift and Variety
- The Semantic Glue (Entity Resolution)
- Bitemporal Auditing
- TPU Readiness

Tier 2 serves as the evaluation ground where the system identifies blind spots. For example, the Tier 2 analytic component compares the SRO's native output against the Tier 1 agent's output, sorting the results into an operational audit matrix:

Scenario	SRO Tool	Tier 1 Agent	Regulatory Classification	Action Item
A	Flagged	Flagged	True Positive	Normal Ingestion
B	Flagged	Missed	Potential False Positive	Send to Tier 3 to inspect SRO over-sensitivity or other nuances causing behavioral shifts
C	Missed	Flagged	Potential False Negative	**Critical SRO Gap** – Enrich with e.g. 13F/ 13H, semantic inference to unveil linked layers

Enrich the Gaps: If Scenario C occurs (the SRO missed it, but Tier 1 Agent caught it), the Tier 2 agent immediately triggers a **Progress Semaphore** lookup,[104] for example: *Who owns this account, the layers of sub-accounts if any? Is it associated with a registered Large Trader (13H LTID)? What are their broader asset holdings according to their latest 13F disclosures?*

[103] https://docs.progress.com/bundle/marklogic-server-understand-concepts-12/page/topics/security.html
[104] https://www.progress.com/semaphore

Semaphore provides the semantic enrichment that grounds **Retrieval-Augmented Generation** (**RAG**).[105] It acts as the semantic enrichment and entity resolution layer. It is an integral part of the Progress Data Platform and sits in front of **MarkLogic**'s multi-model database to analyze complex market anomalies, extract **relevant profiles**,[106] and supply enriched, structured context to the Large Language Model (LLM) or TPU cluster for **deep semantic inference**. Unlike traditional approaches, RAG relies on simple *"keyword matching"* or mathematical vector proximity, which can miss explicit relationships or surface false ones. **Semaphore** instead translates raw identifiers and manages the ontology and taxonomy model, expressing the resolved relationships to **MarkLogic** as RDF (Resource Description Framework) Triples (*Subject-Predicate-Object*).[107] **Semaphore** stitches together siloed pieces of information and draws semantic links such as: [Account X] – (Beneficial Owner) → [Trader Y] – (Controls) → [Algorithm Z].

From an **identity Context** perspective, the Tier 2 agent determines *why* the SRO missed the flag in Scenario C.[108] For example, the SRO may miss a wash trade because the participant used two different broker-dealer accounts that the exchange's internal system could not link. **Progress MarkLogic** semantic identity hub resolves these connections on demand, across siloed records and without a full data reload.

Another use case is addressing the quality assurance (QA) challenge in *port-level settings*. For that, we recommend a **parallel process** to streamline the back-and-forth *"unlinked/ Error"* communications between SROs and Broker-Dealers, and making related resolutions transparent to the SEC. AI agent dynamically wraps mismatch into a lightweight, schema-agnostic Mismatch Notification Document. If the allowable regulatory duration (T+3) is expired without a commonly agreed resolution between the SRO and Broker-dealer, the Tier 2 system would generate a Level 2 Alert prompting the SEC to step-in. See our response to Q21 for an elaborated discussion.

Users are no longer continuously scanning the entire regulatory filing databases against trillions of records. Tier 2 queries the SRO edge nodes via MCP to see what tools and historical data tranches are available for the specific timestamp block. The earlier mentioned Tier 1 MCP Servers retrieve the raw telemetry from **MarkLogic**, package and return it to Tier 2. **Progress Semaphore** takes these multi-SRO responses, normalizes them into unified RDF Triples using the governing ontology. A Tier 2 MCP Server then delivers the aggregated insights over Server-Sent Events (SSE) transport for subsequent analysis in the Tier 3 TPU. Querying the identity graph is **exclusively on-demand** for data that a Tier 1 agent has already determined to be *suspicious activities*. The identified *anomaly* justifies the summons and pulling of relevant profile information in this audit trail from ② and ③, just like the ④ SRO's native surveillance process would.

CAT's outdated architecture as a centralized *honeypot* vault would NEVER be able to do the jobs required here. Consider the complexity of connecting the dots, where the Volcker Rule's covered funds provision[109] may serve as an illustrating example that the CAT is absolutely NO MATCH with **Progress MarkLogic** and **Semaphore**. Besides, there are different echo chambers where firms compete and collaborate simultaneously.[110] "Hunters" type of firms in the value chain upstream versus the "Farmers" type of firms in the downstream are both experiencing consolidation and/or switching gears to focus on DeFi. The value chain smile curve went upside down,[111] and retail orchestrated a *gamma squeeze* of hedge funds. All phenomena attributed to the need of **Semantic Context Awareness** in this Tier 2 *Audit Trail Hub, Case Library, QA*.

[105] https://www.progress.com/blogs/semantic-rag-series-part-4-content-discovery
[106] https://www.youtube.com/watch?v=r02uMclAW_Y&t=707
[107] https://www.progress.com/blogs/semantic-rag-series-part-3-content-preparation
[108] https://investors.progress.com/node/27511/pdf
[109] https://www.databoiler.com/index_htm_files/DataBoiler%20CoveredFund%20ForDiscussion.pdf
[110] https://www.linkedin.com/pulse/rebate-tiering-competitive-pricing-different-market-centers-kelvin-to-u6l2e/
[111] https://www.linkedin.com/pulse/smile-curve-changes-securities-value-chain-evolves-kelvin-to/

More importantly, due to **UNFIXABLE structural deficiencies** – such as perimeter-based security, a lack of element-level encryption, and central administrator compromise risks – the **CAT completely fails to comply** with the Federal **Zero-Trust** mandate under Executive Order 14028; thus, it must be replaced (see our response to Q28 and Q29). The SEC can have a piece of mind with **Progress MarkLogic** and **Semaphore** that force a **Zero-Trust** edge. **MarkLogic** does not just encrypt the database, it secures data at the Document and Element levels. **Semaphore** dynamically classifies and tags content against the governing ontology, and those tags can drive attribute-, element-, or column-level based access control enforced by **MarkLogic**'s Element-Level Security. Because our architectural design ditches the central vault, the data stays distributed across dispersed systems at SROs, Clearing and Settlement firms (DTCC), and the SEC (e.g. EDGAR filings). Attacker CANNOT pivot laterally to steal data of other platforms, *completely eliminating* the **"apocalypse scenario"** inherent to the CAT *honeypot*.

TIER 3 *Reinforcement Learning Loops (**Cross-Market Validation**)*

Upon a verified local anomaly tied to a definitive institutional / individual identity for suspicious activities, the Tier 2 agent passes this enriched data package up to the Google TPU Cloud cluster to perform the final, most complex check.

A modern flash event or systemic liquidity choke is heavily driven by securities inventory management under extreme stress – where clearing limits and margin constraints force market participants to aggressively unload or refuse risk. Also, MEME events can generate sudden, highly concentrated volume spikes that overwhelm traditional market models, breaking native trading venue queues and broker-dealers' internal risk guardrails. When these compounding stress conditions force multiple key market makers simultaneously pull their quotes from the book, it becomes critically important to monitor exactly how these **stress states develop in real time**, and how system alerts volatility interruption mechanisms to do their job in restoring orderly function of markets.

The TPU monitors the speed at which market makers accumulate unbalanced, directional inventory, calculating the exact threshold where margin constraints will force them to aggressively drop their quotes. Simultaneously, the TPU evaluates cross-market correlations to detect if liquidity is vanishing across multiple fragmented venues at the exact same time, identifying a systemic liquidity choke before it triggers a domino effect.

Instead of writing a massive central ledger, the Tier 3 Agent urgently serializes its findings into a lightweight, asynchronous notification, where the TPU works in concert with the Tier 1 agent for *queue saturation monitoring*. The feedback may contain spiking lag signals that indicate a trading venue's native guardrails are beginning to fracture.

The TPU does not just verify occurrence of events; it runs **Meta-Learning Optimization Models** for reinforcement learning. The previously mentioned *"True Positives, False Positives, and False Negatives"* are converted into tensor matrices. When the TPU analyzes the enriched *"SRO Misses"* data blocks streamed up from the Tier 2 **Progress MarkLogic** and **Semaphore** layer, it examines, for example, if a trader is intentionally exploiting structural differences between exchanges, and considers how the SRO may recalibrate its parameters to close those specific vulnerabilities.

In suppressing *False Positives* – where SROs frequently overwhelm regulators with *"noisy"* alerts that represent standard market-making activity wasting compute power and human resources – the TPU confirms these are harmless artifacts of algorithmic liquidity provision. It classifies them as systemic noise, reducing unnecessary footprint.

Consequently, Level 3 Alerts are dedicated to rectifying *False Positives/ Negatives*, as well as escalating the attention of *True Positive* events for prioritize actions – such as informing the Tier 1 agent to scrutinize a target if its next move match a predicted adverse outcome.

Furthermore, the TPU is given specific identity signatures and execution timestamps. It spins up a targeted Matrix/Graph Neural Network processing job to ask, for example: *While this identity was manipulating NASDAQ at 10:14 AM, what were*

their affiliated accounts doing on NYSE, CBOE, or ATS dark pools? Are we seeing the onset signals of a highly distributed cross-market spoofing ring?

Because the TPU cluster is completely blind to 95%+ of the noise generated by the benign market traffic (*Tier 2 Agents extract ONLY the vital few data blocks containing verified onset anomalies and forward ONLY those targeted blocks with bitemporal integrity to the TPU*), its computing requirements drop exponentially. The SEC and/or SRO's cloud infrastructure avoids paying for continuous, massive matrix operations across the entire macro-universe, reserving the premium TPU muscle solely for verifying target packages.

The SEC's regulatory workstation operates strictly as an MCP Host containing an orchestration client. The SEC investigator views and monitor alerts in a compliance dashboard. Operating under a strict **Zero-Trust** perimeter, the Tier 3 MCP Client sends a structured JSON request to the Tier 2 server, where it receives compact, mathematically verifiable cryptographic proofs and high-level metadata summaries rather than raw transaction files. This interaction is executed within a stateful, persistent, and authenticated multi-step conversation session via a secure Server-Sent Events (SSE) transport layer, ensuring real-time regulatory visibility without centralized data hoarding. The raw data stays completely untouched and safely siloed within the SRO perimeters. The host model uses MCP's tool-calling capabilities to dynamically select and trigger specific analytical micro-routines. The TPU and semantic graph layers remain entirely idle until MCP explicitly calls them to process a flagged anomaly. An SRO MCP Server can dynamically reject an SEC tool-call if the digital certificate lacks the necessary **case-specific context or active case validation**. This **Zero-Trust Context Gating** neutralizes the threat of unauthorized mass government *data snooping* or centralized cloud *honeypot* leaks.

Advantages of the NEW

(a) **Protects the SEC and/or SROs from "Alert Fatigue":** Because different SROs use entirely different tech stacks and vendor tools, by handling localized indicators of manipulation at Tier 1, it acts as a universal translator that prevents the central database from flooding regulators with millions of fragmented, disconnected alerts.

(b) **Solves the Multi-Cloud Regulatory Burden:** SRO inputs can be parsed via flexible microservices anywhere, while the heavy semantic tracking sits securely inside **Progress MarkLogic** (keeping the core audit trail database in Tier 2 lean and responsive), and the specialized, deep-learning math runs strictly on **Google**'s cost-optimized TPUs. The Tier 3 loop catches and quantifies system degradation for automated regulatory calibration.

(c) **Matches the Pace of Modern Market Abuse:** Bad actors/ foreign adversaries intentionally spread their manipulative signatures across multiple platforms to avoid detection by any single exchange's internal surveillance. By isolating the identity in Tier 2, this system effectively strip-mines their anonymity before evaluating their global cross-market market footprint in Tier 3. No more T+5 regulatory access (a predatory algo can manipulate volatility bands, mobilize a flash crash, pocket the profits, and alter its code signature days before an SEC analyst even configures the manual SQL query to investigate the anomaly). Our recommended architecture closes this window.

(d) Shifting form hindsight to active prevention:
Traditional CAT surveillance is purely forensic, acting as a post-mortem that explains why a crash happened weeks after the event. By feeding the timely (clearing and settlement) data into an overnight TPU reinforcement learning loop, our recommended architecture becomes an active defense mechanism suggesting update of LULD bands or noise-suppression tokens down to the exchanges to prevent a similar cascade from triggering the following morning.

(e) Structurally reduces the annual CAT cloud burn-rate:

Cloud hosting services cost the CAT project upwards of $159 million annually (though optimized in late budgets down to $77 million – $81 million) because they keep vast tranches of non-anomalous data indexed on high-availability, expensive SSD cloud arrays. Our recommended 3-Tier Distributed Architecture structurally slashes the annual CAT cloud burn-rate by an estimated 70% to 80%, see below for our initial rough financial estimation:

Cloud Budget Category	Legacy Centralized CAT Baseline (Estimated)	3-Tier Distributed Topology (MarkLogic + Semaphore)	Structural Mechanism for Savings
Active Compute (VMs / TPUs)	$45M – $70M	$7M – $12M	Cloud TPUs run strictly on an on-demand burst model triggered only by edge anomalies.
Data-lake Storage (Hot/Warm/Cold)	$50M – $80M	$8M – $15M	Element-level compression; 90% of normal data sits in low-cost MarkLogic cold tiers.
Networking & Egress Fees	$10M – $15M	$5M – $8M	Eliminates central pooling; SROs dedicated, high-security extranets and enterprise cross-connects.
Total Estimated Cloud Burn	$105M – $165M	$20M – $35M	Total Structural Cloud Cost Reduction: approx. 70% - 80%

***