

Dear Elizabeth M. Murphy, Secretary
Securities and Exchange Commission
100 F Street, NE
Washington, DC 20549

Re: Comment on Proposed SEC Rule 617
Reference: SEC File No. S7-11-10

Dear Ms. Murphy,

This comment follows up our conference call on January 20, 2011, with David Hsu, Jennifer Colihan, Rebekah Lin, and Leigh Duffy. Mr Hsu asked that we submit a comment as cited above to proposed SEC Rule 617. The attached material is the requested comment.

We are happy to apply our experience on the SEC Market Oversight Surveillance System (MOSS) project to the CAT project. We believe the operational and technology insights in the comment will enable the Commission to increase CAT regulatory effectiveness, reduce CAT development and operating costs for the Commission, SROs, and market participants, and will integrate the CAT's value across key aspects of the Commission's regulatory programs for markets, issuers, and enforcement.

Respectfully,

Michael Belanger,
President
Jarg Corporation
781-223-0117

Joseph Carrabis,
CRO and Founder
NextStage Evolution
603-577-4575

Wayne Ginion,
Vice President, Enterprise Infrastructure Services
New York Times Company
212-556-4426

David Morf,
Partner, Senior Regional Economics Advisor, and Founding Member
Center for Adaptive Solutions
413-536-0944

Attachment

Market Transaction Awareness: Graph Technology –
The Scalable Framework for Long Term Regulatory Integration
Comment to Proposed SEC Rule 617 – Reference: SEC File No. S7-11-10

Market Transaction Awareness: Graph Technology – The Scalable Framework for Long Term Regulatory Integration

Comment to Proposed SEC Rule 617 – Reference: SEC File No. S7-11-10

Contents

I. Introduction...

- A. Problem Statement
- B. Solution Approach
- C. Implementation and Operation
- D. Costs to SROs and the Commission
- E. Benefits to SROs, the Commission, and the Public

II. Planning...

III. Design...

IV. Implementation and maintenance...

V. Operation...

- A. Overview of Flow to the Commission
- B. Processing Cloud and Open-ended Transparent Flexibility to Expand Cloud
- C. Interpretive Technology and Analytics
- D. Market and Securities Class Extensions
- E. Functional Extension to Optional Global Hourly Credit Rating Task

VI. Governing...

Graphics

- Initial Planning Model
- Graph-based “Abstraction” and Inferencing

Market Transaction Awareness: Graph Technology – The Scalable Framework for Long Term Regulatory Integration

Comment to Proposed SEC Rule 617 – Reference: SEC File No. S7-11-10

I. Introduction...

A. Problem Statement

The Commission's underlying challenge is to identify matters of interest as quickly as possible across all securities classes for which it provides the underlying regulatory context. Related challenges stem from the complex legal-technical-regulatory environment in which the Commission operates. Specifically, market participants are in constant movement to apply a mix of continually evolving trading strategies that leverage structural links across securities classes, as well as external operating and strategy information, to achieve continuing transaction advantages.

The external operating and strategic information can include activity in non-SEC markets, including markets not under US jurisdiction, and global issuer, market, and supply chain events outside the SEC's regulatory purview. Therefore the basic challenge is to track the constant movement of the meaning of trade strategies as the strategies and their meaning continually unfold, interact, and evolve across cross-influencing issuers-issues-markets in the short run and long run.

Securities offerings, market quotes and trades, and clearing and settlement flows, are the partial outcome of strategies and information leverage, not the drivers. As a result, the data collection to oversee the markets needs to start with tracking active strategies and their context, not with initially collecting and tracking the enormous universe of all resulting offer, trade, participant, and clearing-settlement data.

Collecting transaction data without first framing a data interpretation strategy and environmental context leaves the observer behind the constantly moving learning curve that drives the resulting transactions. Also, the work to collect transactions absent their context requires major initial and ongoing investments in money and equipment to capture, store, and assess the huge volume of raw transactions.

The Commission staff is familiar with gathering data into a central location to support analytic activity and case development. However, to pre-define and collect all data prior to interpretation, analysis, and transformation into matters under investigation as appropriate, itself delays the core necessity of first understanding and tracking the constantly moving, related, and cross-influencing information behind the ostensible quote and trade activity.

B. Solution Approach

The proposed solution is to move the strategy intelligence function to the in-market (front end) of the Commission's CAT process, rather than to an external (back-end) CAT process. It casts an intelligent net for tracking key nuggets of strategy-driving information. Its design and operation evolves its regulatory data uptake in ongoing

response to constantly evolving meanings in the global market strategy environment.

C. Implementation and Operation

Operationally, the practical approach is to start with a pilot phase involving a small number of market participants and representative securities classes, learn the ropes at the SRO, market participant, and Commission levels, then scale up the process to include the national market system for which the SEC provides regulatory context. The pilot phase would apply an ongoing plan for small, mixed-staff teams in each SRO (mixed SRO and SEC people). A senior SEC and SRO review and oversight effort at each SRO will be essential to ensure highly qualified personnel are engaged in the hands-on work teams in both the pilot phase and the scaled-up operations.

D. Costs to SROs and the Commission

Compared to the multi-billion dollar effort to take in the universe of transaction data contemplated in traditional audit trail designs, the proposed solution requires orders of magnitude less cash for investment and operation by markets, market participants, and the SEC than estimated by the Commission's initial CAT plan. For example, the budget to accomplish the work amounts to tens of millions of dollars, not billions. At the nine-month mark, the proposed abstracted graph architecture will be operational, adaptively resilient, and sustainable – meeting currently stated CAT requirements – without requiring billions of dollars, nor even hundreds of millions.

E. Benefits to SROs, the Commission, and the Public

The proposed solution by design tracks in-market movements in trading market relationships across SROs, issuers, suppliers, market participants, and their wider context. As a result, the proposed solution remains coherent, relevant, and effective going forward. As a learning engine by design and operation, the solution itself keeps getting more adroit and relevant, not less, the longer it is pursued.

Compared to traditional audit trail designs, the proposed solution is less intrusive on markets and participants. It tracks strategies and their meaning. It learns to pull in data that build out the picture for rapid interpretive analysis and targeted follow-up.

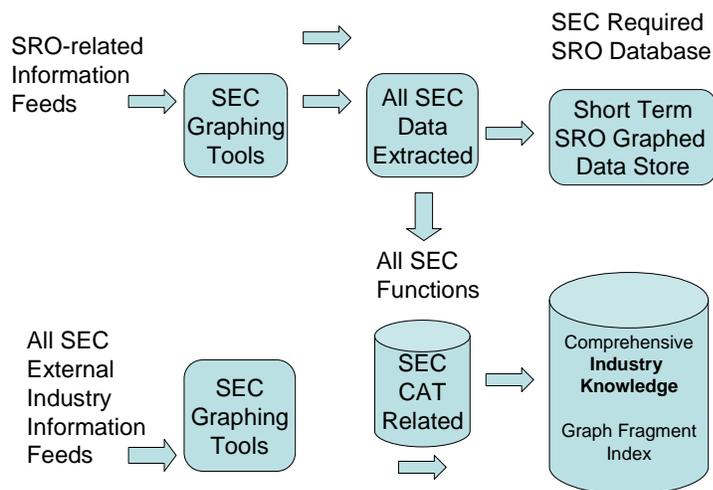
The reduced cost of the proposed solution reduces the public expense involved, both as to SEC budgets for the CAT, and as to SRO and private sector participant outlays for compliance. The increased and continuing effectiveness of the proposed solution enables the Commission and the SROs to protect investors and also to provide more coherent, fair, and economically efficient national securities markets in their US and global context by aligning regulatory technology with best-practice market tools.

II. Planning...

Premise: To observe and interpret events in the securities markets, the Commission needs a platform able to support a continuous intelligence race. The CAT rule's

currently contemplated intelligence operation is equivalent to candles carried by town criers. Operationally, the SEC needs sophisticated and continually adaptive in-market gathering and interpretive tools that apply a design framework as follows...

1. Leave all SRO data in each SRO, exactly where it is being created in each SRO.
2. Install an SEC-controlled software presence within an activities software database to be required within each SRO.
3. Apply existing concept-based graph-matching and indexing technology in each SRO and the SEC. The technology is flexible, scalable, and open to the future integration of all related SRO-equivalent data now subject to regulation by other authorities.
4. The proposed intelligent matching and indexing technology enables the SEC to have secure real-time detection and extraction of SEC-desired activity information from within each SRO's original data bases, and from the Commission's own live information feeds.
5. The scalable approach for achieving the CAT's mission proposed here avoids the replicating, moving, hosting, maintaining, and real-time transaction updating at the SEC of SRO legacy data base content (a major budget and operational saving).
6. Initial Planning Model...



III. Design...

Premise: An architecture and operational model in the Commission's market-based environment needs to ensure that effective regulation can sustain a continual

securities market learning race. For example, a technology platform that applies an enormous legacy “closed” relational data architecture, as currently envisioned for CAT, will allow SROs “technically” to comply with CAT rules. However, relational data architectures are a challenge to modify, which reduces the speed with which the architecture can track a market-based learning curve. The reduced speed of tracking ongoing market evolution will perpetuate continued blindness by SEC regulatory tools to the continually evolving transaction structures and strategies in play within and across SRO and other markets.

1. To level the playing field effectively for current as well as evolving comprehensive market regulation, CAT and its potential extensions need to build upon and apply real-time adaptive graph indexing-based architectures. Applying a graph-based architecture enables the SEC to support an extensible CAT capability to harmonize information awareness for regulators across securities and futures markets.
2. The conceptual graph-based approach being proposed here is fully mature. The technology is used today by millions of people (it underlies Google’s search and other popular applications). It is actively used by hedge funds and by investment bank trading desks. From a technology view, traders currently outpace regulators every day, enabling an increasingly opaque market behavior environment.
3. Absent a direct information awareness extension capability, it would be ironic if a CAT solution optimized for the current “closed” SEC securities data environment had the unintended consequence of deepening the operating framework differences separating futures and securities markets and their siloed regulators, while the money behind the markets continued to game and arbitrage the two market ecologies and other markets with highly fungible cash.
4. The graph technology platform being proposed offers the regulatory user a real-time global alert-based interpretive and actionable view by enabling a virtually infinitely flexible user inquiry framework combining graph-based indexing, search, and software agent technology, with support from intelligent analytic resources.
5. The flexibility and adaptability offered by the proposed system stems from applying graph indexing to the concept cascade represented by a securities market ontology. A market ontology is a tool to nest, control, and apply across SROs the domain knowledge references specific to each SRO (e.g., tire belongs to wheel belongs to car for various cars, or aileron belongs to wing belongs to airplane for various air craft, or marine propeller belongs to marine propulsion system belongs to ship hull for various ships—but all belong to transportation technology).

IV. Implementation and maintenance...

1. The timeframe for the proposed CAT implementation spans some nine months – six months for planning/licensing tools, team training, and processing cloud acquisition, and three months for operational setup at the Commission and in the field.
2. Relevant technology experience indicates that the most productive size for any one

development team, setup team, or maintenance team is four—two SEC staff and two SRO on-site staff. The speed required by the SEC for initial ontology and supporting tools setup at all SROs will drive the number of initial two-person SEC teams in the field at any one time, and thus the system-wide staffing requirement.

3. Operationally, the practical approach is to start with a pilot phase involving a small number of market participants and representative securities classes. The pilot goal is to learn the ropes at the SRO, market participant, and Commission levels. The pilot phase would test the plan for small, mixed SRO-SEC staff teams in each SRO.
4. Following the pilot phase, the two-stage scale up process consists of setup, followed by ongoing operation. The goal is to implement CAT for the national market system for which the SEC provides regulatory context.
5. During the post-pilot field setup stage, an SEC work team (creation team) will visit each SRO. This creation team consists of both SEC and local SRO personnel for an initial four day set up session. The SEC tools and Ontology are populated quickly using existing industry digital libraries. These SEC-SRO team members will know and build into these SRO resident tools both the formal and workday terminology and workflow for each SRO and for the operations of its market participants.
6. Following setup, subsequent ongoing SEC CAT work teams provide SEC-managed mobile software maintenance. The mobile maintenance teams will consist of both SEC and SRO staff, and meet with selected SRO personnel at each SRO location. The two person SEC element in the maintenance teams will visit their SROs once a month to drive periodic tool and ontology upgrading sessions. These SRO-resident (SEC controlled) ontology-graphic tools will be taught the continually evolving jargon unique to each SRO operation, staff, and market participants. The goal is to track the continuously adapting and evolving SRO markets and their participants.
7. After all SRO graph streams to the SEC are operational, the ongoing number of two-person SEC-staffed field oversight teams can be estimated based on each two-person SEC field oversight team visiting or connecting with three SROs per day. Central SEC staff also must be assigned simultaneously to add each SRO-resident ontology intelligence upgrade to the SEC's own internal aggregated ontologies.

V. Operation...

A. Overview of Flow to the Commission

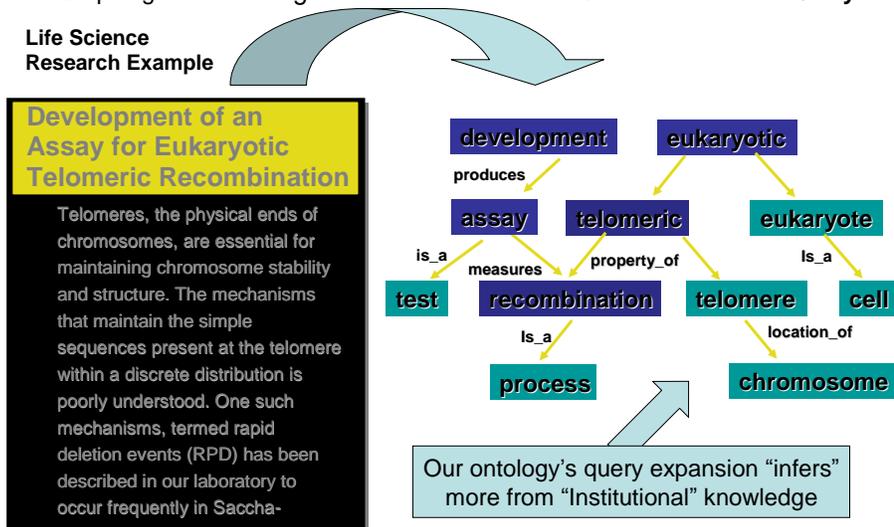
1. At each SRO, a small set of SEC controlled software tools will monitor the appropriate SRO databases and information feeds. These graph-based search and knowledge awareness tools are already available. Tools selected by the SEC for each SRO data location will “understand” that SRO's data formats, all standard SRO terminology, and will be taught any jargon unique to that SRO's staff,

markets, market participants, and financial instruments. These tools incorporate an Ontology, i.e., a nested domain knowledge reference specific to each SRO (e.g., tire belongs to wheel belongs to car for various cars, or aileron belongs to wing belongs to airplane for various air craft, or marine propeller belongs to marine propulsion system belongs to hull for various ships—and all belong to transportation technology).

2. When the SEC software tools are operational, each SRO's information will be conceptually framed (abstracted) in graph format. The “abstraction” is a real-time summary of the meaningful data flowing and residing at each SRO. From that point on, a current summary is continuously and securely flowing to the SEC in a compact, vector-space graph format. The abstraction concept is modeled as follows...

Graph-Based “Abstraction” and Inferencing from Open-Ended Sources

Graphing the meaning of Both The Information Sources and Your Query



3. At the SEC, additional sets of robust software tools will receive the graphs (while retaining their source paths) flowing in from the SROs and participants, and will flow past intelligent “software alerting tools” set up by SEC monitoring staff, and then the fragments will populate a unique, very high performance graph fragment index structure that is searchable and preliminarily interpreted by concept-based analytics (discussed below in “E. Functional Extension to Optional Global Hourly Credit Rating Task”) supporting SEC interpretative staff.

B. Processing Cloud and Open-ended Transparent Flexibility to Expand Cloud

1. The private cloud of SEC-resident CAT servers will be setup mostly to stream the SRO graphs through SEC-resident comprehensive ontology-based intelligence software tools for compliance, alerts, and interpretation. The retained data largely

will be selected by pre-set alerts, and their related information will be sent to the applicable SRO's interpretive team and to SEC oversight and inspection teams, full disclosure units, and enforcement staff. Interpretive tools for SEC use are described below in "E. Functional Extension to Optional Global Hourly Credit Rating Task." Existing SEC tools to assess quote, trade, and news patterns also support SEC staff.

2. The SEC may decide to expose meaningful "abstracts of" related information populated by all the other SROs, without giving any one SRO the capacity to reconstruct any specific SRO-resident data or knowledge sourced from the other SROs. This design could allow each SRO directly to benefit from aggregated CAT data without being able to replicate the internal data of other SROs.
3. Regarding data and report sharing, the SEC will need to distinguish analytics and required data for SEC use only, vs SRO shared access. The SEC should consider not exposing all SEC-based analytics to SRO's as it likely will be difficult to ensure that any analytic activity can effectively "blind" some of the SRO-wide reports or analytics. SROs and market participants will be adept at reading blended CAT data – that's what they do for a living.

C. Interpretive Technology and Analytics

Premise: To follow behavior and money across transactions, participants, and issuers, SEC social and media interpretative analytics and media mining should be applied to related money-based transaction activities, topics, and data sources. Key examples:

1. Quarterly corporate conference calls, MD&A disclosures, financial disclosure and footnote data, and the recommended hourly credit rating flow described below in section "E. Functional Extension to Optional Global Hourly Credit Rating Task;"
2. Commodities, Currencies, Banking, Proprietary Information, Intellectual Property, Black Pool interpretive views, Communication Advantage Regulations, Global Regulatory Reporting Standards Normalization, Intelligence Community funds flow tracks; and
3. Corporate filings in state and federal courts and agencies—including submissions to federal appellate, district, and bankruptcy courts—for antitrust, litigation, intellectual property, and corporate actions.

Any interpretive technology solution to a social challenge needs to have a working understanding of human behavior and motivations in order to be successful. More importantly, the technology solution must not itself demonstrate the behaviors and motivations it monitors and reports on, in order not to elicit social polarity reactions in response to the interpretive technology.

Therefore the need exists for a technology that recognizes, understands, and reports on human behavior, motivations, and emotions while not itself displaying those traits.

The interpretive technology should work on two levels – observation and review.

The first level for the purposes outlined here is to observe the online community being regulated. This observation should take the form of an “anthropologic blind.” This means observing community members when they are “being themselves” as opposed to when they are in a controlled situation (focus group, marketing study, survey instrument, etc.). Controlled – or *synthetic* – situations can only demonstrate synthetic results. Thus the observational aspect, the anthropologic blind, needs to be something that already exists and is routinely ignored by members of online communities—examples include embedded traffic analysis technology or embedded graph flow technology.

The second level of technology is the actual analytics review of the data traffic. The purpose of this review is to determine if there are agendas, intents, or patterns in the traffic that need to be flagged as possible matters of market oversight interest.

Operationally, the traffic interpretation technology assesses arcs of issuer and market participant behavior by linking the interpretive tools to the graph flows. Correlating precipitating activity, agendas, intents, and patterns with subsequent market behavior and full disclosure compliance activity, especially if linked with known tools to assess quote, trade, and news patterns, offers direct oversight value to both the market regulation and full disclosure programs.

The relevant scope of precipitating activity would first link interpretive technology to the initial CAT universe, augmented by related money-based transaction activities, topics, and data sources cited above, especially court filings, quarterly corporate conference calls, MD&A disclosures, financial disclosure and footnote data, and the credit rating flows recommended below in section “E. Functional Extension to Optional Global Hourly Credit Rating Task.” The non-intrusive and reality-based interpretive analysis of continually evolving (time-tagged) intent contained in the fragment flows and other data sources drives increasing returns to scale for the CAT.

Existing technology patented by NextStage Evolution, LLC, has the capacity to perform the observational and review levels of interpretive analysis in real time as individuals go about their normal business in the online SRO, market participant, and issuer communities. This technology has been in use in the business world since 2001.

D. Market and Securities Class Extensions

Premise: Appropriate regulators across all global financial activity also can field teams to setup and populate graph-based search tools and ontologies at source locations to extend the graph framework already discussed.

1. From a technical perspective, the same high performance graph fragment index structure set up for the initial CAT universe, for which the SEC provides regulatory context, also can absorb and interrogate real-time output flowing from any other added markets. These could include futures, black pools, other traded instruments, and surrounding environments as noted below in section “E. Functional Extension to Optional Global Hourly Credit Rating Task,” without

reducing system performance – for example, regarding futures, to include all related derivative and risk management instruments while maintaining performance, just add servers to the SEC’s cloud and apply to the futures markets the implementation and operation phases outlined in Part IV above for securities markets across underlying securities and derivatives.

2. Search indices can serve their alerts to each added market’s own market monitoring teams. For coherence, one top-level CAT index (with cloned backups) should provide the expanded real-time summary view to SEC market oversight staff and to SEC interpretive technology (or to similar staff and technology at systemic oversight entities created pursuant to Dodd-Franks) for all the financial behavior activity that is being monitored.
3. For regulatory harmonization, the integrated regulatory graph-based information index operation includes participants and securities for all related derivative and risk management instruments. Much like managers monitoring the real-time performance of a very large data network such as the Internet, market regulators must build upon a graph index architecture. This architecture is necessary to command an integrated real time view, in any required combination, over any, several, or all of the following critical, yet currently siloed, financial market operating components, regarding participants and securities for all related derivative and risk management instruments, such as...
 - Market participants
 - Real time market transaction monitoring
 - Clearing
 - Settlement
 - Real time green-yellow-red global issuer activity and risk alerts
 - All credit ratings of securities updated (optional project) every hour.
 - Alert drill-downs for market participant traffic activity
 - Intent interpretation
 - Corporate governance
 - Executive compensation

E. Functional Extension to Optional Global Hourly Credit Rating Task

Premise: It is strongly recommended that the Commission leverage CAT effectiveness by linking market and issuer oversight in real time with hourly global credit ratings. The impact transforms CAT into a strategic multi-program SEC asset that cross-leverages three core Commission programs – full disclosure, market regulation, and enforcement. A strategic (multi-program) regulatory asset offers *increasing returns to scale* and is in a stronger regulatory and budgetary position than a more limited (single program) asset.

The credit rating work involves an optional-cost task to track and assess global traffic flows for issues and issuers. The task work represents the equivalent development effort and operational overhead of three more two-person SEC-staffed CAT field teams.

The credit rating process would generate globally-aware hourly credit ratings of securities using graph-based index and interpretive tools. The direct output is an hourly report of credit ratings (green-yellow-red flags with drill-down to supporting flag detail). The Commission-wide, cross-program operational and public benefits are two-fold...

1. The full disclosure program's public benefit is to support the full disclosure program's selective review process by adding market intelligence regarding issuer behavior relative to market events and market strategy patterns to the Commission's selective review criteria and to the assessment of full disclosure compliance in selected filings;
2. The market regulation and enforcement program public benefits are to enable patterns of issuer activity to illuminate and accelerate the interpretation of market events and market strategy patterns in a CAT market regulation and enforcement framework that must allocate and target scarce resources to identify, prioritize, and pursue coordinated matters under inquiry.

VI. Governing...

Premise: Behavioral incentives and intermarket trading strategies combine to offer major challenges to effective, timely, and affordable oversight. The transaction context and global ecology for securities market oversight includes accelerating scope, velocity, and leverage due to new financial instruments, processing tools, and communications technology. The transaction context can cause regulatory efforts aimed at investor protection, systemic stability, and economic flexibility to face challenges in locating and resolving destabilizing behavior if the regulatory efforts lack the tools to track the results stemming from today's market incentives, technology resources, participant business structures, and global economic environment. Both SRO and SEC regulators share a need for a new market oversight leverage point.

1. From the SEC's experience in prior efforts to exercise market oversight and surveillance, it's clear that any viable CAT process must be designed to evolve at low cost the ability continuously to stay in focus as markets-participants-issuers evolve their offerings and behaviors in response to their own economic ecologies. Therefore, it's a critical design element that the CAT search tools interpret market and issuer actions. The interpretive capability enables continuously relevant, complete, and coherent regulatory learning, oversight, and investigative pursuit. An interpretive CAT capability ensures that SEC, SRO, and potentially other market regulators, can provide operational continuity, coherence, and completeness year over year, without the SEC or SROs annually spending billions.
2. It's important to consider the SEC's ability to address SRO concerns about CAT. Specifically, the ability of CAT search tools to act as a continuously evolving interpretive inference engine adds to the ability of the SEC (and of regulators such as the CFTC and the SROs themselves) to assess the scope of the data that must be tracked to accomplish the missions of investor protection, market integrity, customer and intermediary protection, and fungibility and competition across securities products and markets, in timely context with the actions of securities

issuers.

3. One responsive outcome of the interpretive capability of the recommended CAT search tool approach is to offer relief to SROs concerned by the scope and cost of direct data capture as contained in the draft CAT rule. Specifically, opportunity one is the potential for interpretive and inference capability of CAT search tool technology to reduce the routine CAT data collection scope. The second key value to SROs and the SEC is to leverage the CAT's interpretive and inference capabilities in order to achieve more efficient and effective monitoring, and to achieve a more rapid and refined pursuit by SRO and Commission staff of ensuing matters under inquiry.
4. In an interpretive search tool environment, the core governance activity in practice will be to staff and sustain a senior SEC and SRO review and oversight function at each SRO to ensure that highly qualified personnel are engaged and productive in the hands-on work teams in the pilot phase and full-scale operations. Data collection boundary questions inherent in tracking market evolution need to be addressed at the SEC-SRO oversight level, with policy questions elevated to existing SEC-SRO venues for timely resolution. The goal is to sustain the CAT's ability to serve the public, the SROs, and the Commission.

Respectfully,

M. Belanger, J. Carrabis, W. Ginion, D. Morf