U.S. SECURITIES AND EXCHANGE COMMISSION

ROUNDTABLE ON TECHNOLOGY AND TRADING: PROMOTING STABILITY IN TODAY'S MARKETS

Tuesday, October 2, 2012
10:00 a.m.

U.S. Securities and Exchange Commission
100 F Street, N.E.
Washington, D.C.

SEC COMMISSIONERS PRESENT:
Hon. Mary Schapiro, Chairman
Daniel Gallagher, Commissioner
Troy Paredes, Commissioner
Elisse Walter, Commissioner

PARTICIPANTS PRESENT:
Sudhanshu Arya
Gregg Berman
Jim Burns
Thomas Bayer
Robert Cook
Tom Eady
Amy Edwards
Robert Fishman
Chris Isaacson
Andrei Kirilenko
Amar Kuchinad
Dave Lauer
Nancy Leveson
Craig Lewis
Lynne Markus
Jamil Nazarali
Lou Pastina
Dawn Patterson

PARTICIPANTS PRESENT (Continued):

Chris Rigg
Jon Ross
Todd Scharf
David Shillman
Heather Seidel

C O N T E N T S

P R O C E E D I N G S

MR. R. COOK: Good morning, everyone. My name is Robert Cook. I am the Director of the Division of

Trading and Markets at the SEC, and it's my pleasure to welcome all of you here today to the Commission's Technology and Trading Roundtable.

We're pleased you are all able to join us today for what we anticipate will be an informative discussion. I think we're going to be bringing in some more chairs in a moment for those of you standing in the back. So bear with us, but we're going to get the roundtable rolling in the meantime.

First, let me just give a quick overview of the day. We'll have two panel discussions. This morning's panel will focus on error prevention. We will hear from technology experts about best practices and practical constraints for creating, deploying and operating mission critical systems, including those used to automatically generate and route orders, match trades, confirm transactions, and disseminate data.

We expect to break from this panel at about 12:15. We will then pick up again at around 2:00 p.m. for our afternoon panel, which will focus on error response. We expect to learn from our technology experts about how the markets might address erroneous quoting and

0006

trading activity when it does occur in order to limit its impact, through independent filters, objective tests and other real time processes and crisis management procedures to detect, limit and possibly terminate the erroneous activity.

Before we begin with our first panel, it's my honor to introduce Chairman Mary Schapiro and invite her to offer some opening remarks.

CHAIRMAN SCHAPIRO: Thanks very much, Robert. I want to thank you and your team in Trading and Markets and in RiskFin for all the great work you've done organizing today's roundtable.

Good morning, everyone, and I also want to thank all of our panelists for taking time to share your thoughts with us on market technology. And I also want to thank all of those who have already written in with comments. You've given us a number of very thoughtful recommendations that will become part of the record of what we're doing here today.

To an extraordinary extent, the stability of our securities markets is tied to the technological infrastructure of those markets. As with virtually every industry, technology brings many benefits, and in that respect our markets are no different. Thanks to technology, our securities markets are more efficient and accessible than

0007

ever before.

But we also know the technology has pitfalls, and when it doesn't work quite right, the consequences can be severe. Just imagine what can happen if an automated traffic light flashes green rather than red, if

a wing flap on a plane goes up rather than down, if a railroad track switches and sends the train right rather than left.

Similarly, there could be significant consequences for technological errors in our markets as well. Trading can be disrupted. Investors can suffer financial loss. Friends can be imperiled, and confidence in our markets broadly can erode.

Today's roundtable will help us think through the issues and the steps we need to take to ensure that our markets remain the most robust, efficient and stable in the world.

There are two basic concerns we need to focus on that are highly interrelated. These are, first, the structure of our markets, such as multiple execution venues, the presence of high frequency trading, dark pools and the like; and, second, the infrastructure of our markets, as in the technology that undergirds trading activity.

To provide some perspective, in January of

0008

2010, I asked the staff to begin a comprehensive review of the equity market structure. It was a review that included gathering views on everything from the impact of high frequency trading to the continued rise of dark pools, to the complexity of a multi-venue market system. The focus was not so much on the infrastructure of our markets, but on the way the markets and market participants operate and behave.

Four months later when disorderly trading activities in the S&P E-Mini Markets spread to the equities market causing what is now known as the Flash Crash, we as an agency were well positioned to respond. Working with the exchanges, we quickly put in place a series of measures that have since helped to reduce the likelihood of another event like that from occurring. Within days we summoned the heads of every exchange to the SEC to hammer out common sense approaches to bolster our markets. And as a result of our efforts, we now have in place single stock circuit breakers to prevent stocks from falling too far too fast, and we have approved a more advanced limit up/limit down mechanism to limit excessive volatility.

We now have in place a ban on stub quotes and rules clearly defining when a trade can be broken so as to help avoid circumstances that can lead to disorderly

0009

trading. We now have in place rules banning naked access and requiring rigorous pre-trade risk controls designed to help mitigate disruptive trading at the source. And we now have rules requiring large traders, many of whom use high frequency trading strategies, to identify themselves so that the Commission can better monitor and analyze their trades, a process that other regulators

overseas are beginning to emulate.

Additionally, and perhaps most importantly, we have adopted a rule that requires SROs to develop plans for the first ever consolidated audit trail, a feature that will allow regulators to surveil and reconstruct trading across many platforms.

But there are issues around market structure and the conduct of market participants that we should further examine, including the high volume of cancellations, a proliferation of order types, transparency, high frequency trading generally, potentially manipulative trading strategies, and data latencies for public investors, to name just a few. These issues will require attention, and we are committed to addressing them.

Today's roundtable will focus more specifically on infrastructure not only because of its importance, but also because I do worry that this issue is at risk of

0010

being lost and subsumed by the broader debates regarding market structure. After all, issues that get lost often do not get resolved, and these matters of infrastructure are essential to any holistic approach to improving how our markets operate.

Consider for a moment the IPO of BATS on its own exchange and the IPO of Facebook on the NASDAQ exchange. Now, there are many views regarding the fragmented nature of simultaneous trading across multiple venues. I believe these IPO events evidence a very different set of concerns. Both events involved one of the few single exchange processes that remain in an otherwise fragmented market, namely, building a single order book and crossing trades at a single price to open trading for a newly public company.

In the case of BATS, it was a flaw in new software code designed to conduct a corporate IPO auction. That mistake caused the matching engine for tickers in a certain range to enter into an infinite loop making these tickers, which included the symbol for BATS itself, inaccessible on BATS.

In the case of NASDAQ, the IPO software was designed to accept cancellations submitted while the final IPO price, or the cross, is being calculated. Cancellations received during this time changed the order

0011

book. By design the system recalculated the final IPO price to factor in the new state of the book, but again, changes were received before the system could print the opening trade, which resulted in additional recalculations. This condition persisted resulting in further delay of the opening print.

These single exchange problems are not a result of complexity or fragmented markets, but rather a result of more basic Technology 101 issues.

Consider as well the events this summer with Knight Capital, a trading firm that had just installed trading software that was intended to send orders to the NYSE's new retail liquidity program. Instead the software wound up sending a ton of orders into the market. As the market data that morning revealed, the software did not create patterns of rapid orders and cancels. Rather, the data showed a massive amount of orders resulting in executed trades that caused Knight Capital to accumulate significant and unwanted positions.

This type of problem, as with the IPO mishaps, was again the result of basic Technology 101 issues. Events like these demonstrate the core infrastructure and technology issues that can be problematic in any market structure. However, though for today we are focusing on infrastructure issues, it is important to recognize how

0012

the overall structure of our markets can affect how our infrastructure is designed and implemented.

For example, we have a very competitive market environment in which rapid innovation and speed to market may compete with diligent testing and validation of the technologies that support such innovation. Our multi-venue, interlinked market structure also means that an infrastructure failure by one part or one venue may cascade into other venues and affect many other parties.

And, of course, the inherent speed of trading, which itself is partly a result of the competitive nature of our markets, means that even small, short-lived infrastructure issues can cause drastic harm.

To be sure, several of the measures we have already approved have helped to strengthen our markets, even in the face of potential and inevitable technological errors. Indeed, several of the post Flash Crash reforms, such as the revisions to the clearly erroneous rules, helped limit the impact of the Knight Capital episode on other market participants.

But limiting any harm resulting from technological errors is not as good as preventing the error in the first place, which is why we have instituted clear rules that require firms with access to our markets to have controls in place to reduce the chance of such

0013

errors.

But perhaps the strongest message from the Knight Capital episode is that the party committing an error may very well end up bearing a massive financial loss, and if there is a financial loss to be incurred, it is the firm committing the error that should suffer that loss, not its customers or other investors. That more than anything sends a wake-up call to the entire industry.

Nonetheless, our concern is not whether a single firm might fail, but whether it causes collateral

damage to investors and their confidence in the integrity and stability of our markets. So I'm very pleased the industry has been working overtime in the aftermath of the Knight Capital episode to address these issues, and I'm pleased our roundtable has spurred discussion.

By focusing on the underlying nature of these incidents and hearing as we are today from experts in technology, I hope we can address these issues in an efficient, effective and expeditious manner.

I will now turn the chair back to Robert Cook, Director of the Division of Trading and Markets, and his staff who will serve as moderators of today's discussion.

Thank you.

MR. R. COOK: Thank you, Chairman Schapiro.

0014

Our panel for the morning session is entitled "Preventing Errors through Robust System Design, Deployment, and Operation." Before beginning this discussion, however, I want to introduce who we have around our square table here.

First and foremost, it's my honor to acknowledge with appreciation that in addition to Chairman Schapiro, we're delighted to be joined by Commissioner Walter, Commissioner Paredes, and Commissioner Gallagher.

In addition, there are a number of representatives of the Commission staff here who will be asking questions of our panelists. From the Division of Trading and Markets, Jim Burns and Gregg Berman will be assisting me in moderating the panel, and we're also joined by David Shillman, Heather Seidel, Todd Scharf, Amar Kuchinad, and Tom Eady.

I should just note that Todd Scharf is acting head of our Automation Review Program, or ARP Program, and he's also the Commission's Chief Information Security Officer.

From the Office of Information Technology we have Tom Bayer, who is the Commission's Chief Information Officer, and from the Division of Risk Strategy and Financial Innovation, we have Craig Lewis, the

0015

Commission's Chief Economist, and Amy Edwards. And from our Office of Compliance, Inspections and Examinations we have Robert Fishman this morning and I believe we'll have Dawn Patterson this afternoon.

I should note for the record that any questions or observations that any of us on the staff might put forward today reflect our own views and do not necessarily reflect the views of the Commission, any Commissioner or our colleagues on the Commission staff.

Of course, the folks we're all looking forward to hearing from the most are our distinguished panelists who we're delighted to have with us here today. They were invited to speak today because of their

expertise and practical experience with the design, deployment and operation of complex technology systems. We very much appreciate their willingness to travel here and to share their observations and insights with us.

I'm going to ask each of the panelists to briefly introduce themselves in a moment, but first I'd particularly like to extend a special welcome to two of our guests who come from academia because we've asked them to provide a brief presentation at the beginning of each of our panels to provide some broader context for the discussion.

Sitting to my left and part of our first panel

0016

is Dr. Nancy Leveson. Dr. Leveson is Professor of Aeronautics and Astronautics and also Professor of Engineering Systems at MIT. She is an elected member of the National Academy of Engineering. Professor Leveson conducts research on the topics of system safety, software safety, software and system engineering, and human-computer interaction.

In 1999, Dr. Leveson received the ACM Allen Newell Award for outstanding computer science research, and in 1995, the AIA Information Systems Award for developing the field of software safety and for promoting responsible software and system engineering practices where life and property are at stake.

Professor Leveson has kindly offered to provide the opening remarks for our first panel.

And sitting to my right is Dr. Lynne Markus. Dr. Markus is the John W. Produska, Sr. Professor of Information and Process Management at Bentley University. She's also a research affiliate of MIT Sloan's Center for Information Systems Research. Professor Markus' teaching research and consulting interests include enterprise and interorganizational information systems, the unintended consequences of information technology and risk management strategies, and IT in the mortgage industry.

She was named Fellow of the Association for

0017

Information Systems in 2004 and received the AIS Leo Award for exceptional lifetime achievement in information systems in 2008.

Professor Markus has kindly offered to participate on and provide opening remarks for our afternoon panel, but for this morning's panel she is sitting here with us to ask questions and provide observations as part of the discussion for the first panel.

So I'd like to ask the rest of our panelists to briefly introduce themselves and the organizations that they represent, after which I'll ask Dr. Leveson to start us off with her prospective. So why don't we just go right on down the panel and people can introduce themselves?

MR. ARYA: Good morning, Chairman Schapiro, Commissioners and division staff, and thank you for the opportunity to participate in today's discussion of technology in trading. My name is Sudhanshu Arya, and I'm a Managing Director at ITG. ITG is a global broker that executes as agent on behalf of institutional investors and broker-dealers both on and off exchange. We represent approximately three percent of daily volume in U.S. equities.

At ITG I'm responsible for the full technology

0018

cycle for liquidity management business. This includes all of our arithmetic trading, routing to exchanges of other trading venues, market office functions, and Posit, one of the largest, most established dark pools in the world.

In our search for best execution, we're required to tackle a number of technical challenges, including track and climb positions, intermediating order routes in parent-child relationships, managing open field and canceled orders in the marketplace, and tackling latency, throughout and compliance demands associated with operating in today's environment.

Our aspiration for error prevention is simple. We strive for zero errors in production. Although this goal may seem unrealistic, it establishes a culture of prevention and planning. An error free production environment is key to our license to innovate. In an extremely competitive environment we are highly motivated to minimize trading risk and errors.

ITG spends significant resources on prevention and recovery mechanisms that go above and beyond regulatory requirements. Today's roundtable discussion represents a good opportunity to enhance industry-wide efforts to prevent errors, but errors happen and always will. When they do occur, efficient recovery is

0019

critical. Industry coordination based upon agreed upon metrics and established communications plans and protocols are crucial.

We look forward to contributing to the industry's efforts to establish trading technology best practices, and we commend the division for organizing today's event.

Thank you.

MR. ISAACSON. Good morning. I'd like to thank Chairman Schapiro, fellow Commissioners and the SEC for arranging this roundtable on market technology and inviting me to participate.

I'm a Chief Operating Officer of BATS Global Markets, a founding employee, and was an original developer of the BATS trading system. I oversee technology and operations at BATS.

During my tenure at BATS, we've launched an ECN

to the National Securities Exchanges, a European multilateral trading facility, a U.S. Options Exchange, and completed the acquisition of Chi-X Europe, and successfully upgraded the Chi-X Europe to the BATS technology.

With the BATS team, I also responded to the unfortunate technical error that caused the withdrawal of our own IPO and have helped manage through the resulting

0020

crisis.

Financial markets, especially the U.S. equity market, have changed dramatically in the last 15 years. Numerous changes in regulation, such as the newer handling rules, Reg. ATS and Reg. NMS, along with mass adoption of electronic trading, have permanently transformed the markets. Well publicized events like May 6th and recent technology problems have fueled concerns about the quality of market structure and overall market confidence.

The technology problems this year have also raised concern about the level of control surrounding technology that drives our capital markets. While realizing the complex systems inevitably fail at times, today I look forward to discussing how we can better prevent these failures and minimize their impact on the broader market when they do occur.

Regarding the prevention of errors, I'd like to discuss three key ideas today. First, better policies and procedures to ensure adequate time is devoted by market centers and participants to testing in different environment before production rollout.

Second, the requirement that all market centers support test symbols that allow for testing of new trading strategies and order types in a live production

0021

environment.

Thirdly, more tools developed for and greater involvement of business and compliance personnel to enhance coverage of possible test case scenarios.

Regarding the responding to errors and minimizing their impact, I suggest the following idea:

Document the practice of remediation of known potential crisis scenarios;

Enhance the monitoring to expedite investigation of unforeseen issues and their root causes.

A primary example of enhanced monitoring should be the receipt of real time execution drop copies from market centers by participants in order to reconcile in real time the participant's record of executions versus the market center's record of executions.

Thirdly, know who your key stakeholders are and be prepared as part of your crisis planning to reach out to them and keep them informed in as near real time as possible.

And fourthly, and finally, kill switches at the exchanges ideally invoked in a coordinated fashion across the industry based on centralized, real time position information reported to the DTCC.

I look forward to discussing many of these ideas among others with my fellow panelists and the SEC today.

MR. LAUER: Good morning, Chairman Schapiro, Commissioners and the SEC Division of Trading and Markets. Thank you very much for hosting this roundtable on such a critical issue, and thank you for having me. My name is Dave Lauer, and I'm currently consulting for Better Markets.

At Better Markets we advocate reforms to help curb the abuses in the marketplace. I primarily focus on the abuses of unfettered technology, high frequency trading, and predatory trading practices.

As well I'm consulting for IEX Group, which is a private company building an investor-owned and investor-focused equity market center.

Since 2005, I have been studying the electronic markets, first, while building low latency trading infrastructure which was measured in microseconds, building actual hardware devices. That was at a start-up where we worked with most of the top high frequency trading desks and algorithmic execution desks.

I then took a job after that as a quantitative analyst in charge of both research and trading both high and lower frequency trading strategies.

After working for a time focused on both quantitative research as well as back-testing and trading these strategies, I got substantial experience in understanding the complexities that high frequency trading strategies face in terms of order routing and the speed at which they had to move and make decisions.

After witnessing the Flash Crash and many other smaller incidents on a near daily basis, I decided to leave the industry and try to use what I have learned to help try and fix some of the problems that we've all witnessed.

I think many of us on the panel agree on some of the issues, such as enhanced testing and quality assurance are necessary. While many of the more mature firms in the industry do this, many of the less mature ones do not and often act with very much a technology start-up mentality.

I believe technology has moved so quickly that most market participants have not been able to keep up. Maybe internally they have and they stay at the cutting edge, but very few firms at the cutting edge, high frequency trading firms included, understand the true complexities involved when all of these algorithms are

interacting with each other in the marketplace and the nonlinear incidents that happen, such as the Flash Crash.

I would urge the SEC to recognize that solving the problems of market structure, fairness and

0024

orderliness or now technology issues as the technology is the driving force of the markets, and I would also urge the SEC to consider the tenets of the most disruptive technology force of our time, the Internet, which is based on openness, transparency, and a departure from the proprietary ways that have marked the technology revolution on Wall Street up until this point.

I believe that with a true open, market-wide surveillance system we can confront many of the issues that we see by using and leveraging more advanced technology algorithms.

Thank you.

Good morning. I'd like to begin by extending my thanks to Chairman Schapiro, the Commission and staff for organizing this roundtable.

My name is Jamil Nazarali, and I'm Senior Managing Director at Citadel Securities, a leading retail market maker in equities and listed equity options.

I appreciate the opportunity to participate in this roundtable and to help advance the dialogue to ensure market stability and integrity in today's increasingly automated trading environment.

I joined Citadel over a year ago and have been active in the industry throughout my career, most recently spending ten years at Knight Capital Group,

0025

where I was responsible for electronic trading.

Citadel Securities trades approximately 13 percent of U.S. consolidated volume in equities and 21 percent of U.S. listed equities volume. We operate an industry leading market making franchise and an institutional markets platform. Our team serves a diverse client base and is a major liquidity provider to retail investors in the U.S. and around the world.

From these vantage points I've been able to see how for years automated trading systems have provided enormous benefits for everyday investors, dramatically lowering trading costs, improving market transparency, and increasing market efficiency.

At Citadel, we have a vested interest in sound and stable markets. We are hopeful that today's discussion and follow-on work will help to ensure that investors continue to benefit from automated trading.

Thank you.

MR. PASTINA: Good morning. I, too, would like to thank Chairman Schapiro, the SEC Commissioners and the staff for putting together today's roundtable to discuss some of the most pressing issues facing the U.S. capital markets today.

My name is Lou Pastina, and I currently manage the NYSE Cash Equity Market Operations Group, including

0026

the development schedule, the tactical operating plans, the NYSE trading floor, and the electronic trading component of the NYSE equity market.

For the past 30 years, I have been involved in some capacity in nearly every aspect of the exchange's business. During this time frame, I've witnessed the evolution of the execution process in equity markets from the use of pencil and paper on the trading floor of the NYSE and NASDAQ's electronic quote screen phone system to a sophisticated web of servers that include several tiers of fully automated execution services via broker-dealer internalization, ATSes, ECNs, and exchanges.

This web of technology has allowed for an unprecedented level of investor interaction that could, if fully realized, have a tremendously positive outcome for all investors. A component of maximizing this technology requires market participants providing execution services to implement a greater level of oversight of their development and implementation of these services which will provide consumers with a greater level of confidence in the U.S. equity market infrastructure.

As you know, self-regulatory organizations operate under the direct purview of the Securities and Exchange Commission and participate in the SEC's

0027

automated review policy program, as it pertains to the review of trading technology and security. The ARP Program conducts on-site annual reviews, quarterly updates calls, and engages in regular dialogue with us on matters relating to the technology development, quality assurance, and security functions of our exchanges.

Although the program is voluntary, we have always treated it as mandatory and generally believe it serves several purposes, including assisting regulators in identifying issues that might be occurring across several SROs.

As the industry and the SEC consider solutions that will enhance the stability of the capital markets, we believe it is important for any standards to be uniform across the industry and any additional requirements that are instituted should not exclusively place additional requirements on exchanges.

Given the rapid rise in the number of trade execution venues and the technology used by traders to access each of those venues, it remains in the interest of all market participants to universalize those standards. Given the topic of this panel is preventing errors through robust system design, deployment and operation, there are two primary points I hope to outline today.

0028

The first is the establishment of common standards that can be used in the industry for technology deployment.

The second is a possible technology solution we believe will address a volume driven event similar to the one that occurred at Knight Securities.

Everyone involved in the securities transaction business recognizes that there is no single way to go about delivering innovative technology options and solutions to market participants. This includes the differing functionalities available at venues as well as the deployment and testing of new technology introduced into the trading environment.

However, given the complexity and interconnectedness of trading technologies that exist today, we believe the industry should consider establishing a common set of standards or best practices for all execution venues to follow regarding the implementation and deployment of trading technology.

In particular, we believe that the industry could benefit from common standards for system development life cycles requiring minimum standards for specifications, coding, testing and implementation; employing capability models and certifying to minimum levels, perhaps making those metrics public; implementing

0029

standard testing requirements for large scale deployments, including unit regression integration, customer-industry recovery and capacity testing, and having capacity and recovery standards.

Consideration should be given to establishing minimum capacity standards and recovery standards for all execution venues.

Establishing a common set of standards that are broadly accepted can go a long way to minimizing the confusion imbedded in an already cumbersome set of processes designed to eliminate problems from occurring before they're placed into production.

Most of the existing market center controls today are focused on volatility in the market rather than volume, as was the case with the most recent major market technology issue. To that end the SEC implemented 15c3-5, market access control rules from July 2011, which are designed to require broker-dealers to establish practices that should prevent market-wide issues as a result of excessive erroneous volumes.

While it has yet to be determined if these rules are being implemented with the benefit they seek to provide, we believe that any additional rules that technology solutions discussed would be an added layer of protection for the market and not in lieu of 15c3-5.

0030

As I mentioned, there are already several

methodologies exchanges use to curtail trading disruptions. At the NYSE, we have three primary controls: throttling, liquidity replenishment points, and line monitoring.

In addition, there are market-wide methodologies which include single stock circuit breakers which will be replaced with limit up and limit down; market-wide circuit breakers; and clearly erroneous execution rules. Each of these items has worked as patches to address the problems they are designed to affect. However, none of them focus on excessive volumes from a single source.

As you may know, subsequent to the issues at Knight, NYSE Euronext helped organize a working group of market participants from nearly every corner of the trading universe, including other SROs, broker-dealers, market makers, proprietary traders, and buy-side firms. One of the premises of the discussion was that although 15c3 is designed to establish controls at the individual broker-dealer level, it does not protect the industry as a whole if the technology problem occurs at a broker-dealer regardless of their intent to comply with the rule.

The working group recently submitted a comment

0031

letter outlining a possible kill switch solution that we believe will provide a second layer of security should a volume based disruption occur again. As outlined in more detail in the letter, we believe a kill switch could be designed at the exchange level that would halt quoting and trading activity of a broker-dealer if it exceeded the established peak net notional volume threshold set by the broker-dealer.

We believe that the broker-dealers would be best suited to choose what their peak net notional volume threshold should be within reasonable measures based on prior trading volumes.

In addition, the exchange could include functionality that would automatically send an electronic message to the broker-dealer should it come within a certain percentage of its threshold, and unless the threshold is increased the broker-dealer's access to the change would halt upon reaching the threshold.

While there are other specifics around such a proposal that must still be discussed, we believe that a kill switch-like solution may be the most widely accepted solution that can be implemented on a reasonably short time frame.

In addition, we believe that there are several market participants that may have longer term solutions

0032

to these issues as well.

Thank you, again, for the opportunity to participate in today's discussions. We intend to

continue our efforts to work with our colleagues in the industry to come to an appropriately measured solution and look forward to answering any questions.

MR. RIGG: Good morning. First of all, I would like to thank Chairman Schapiro and the Commissioners and the staff. I very much appreciate the opportunity to participate in this panel.

My name is Christopher Rigg. I'm from IBM's Global Business Services, which is our consulting business unit, and I lead our Application and Innovation Services Group focused on the banking and financial markets practice.

That group is essentially the organization within IBM that focuses on helping our clients implement either custom develop or integrate custom, large, complex systems integration, and I believe I can contribute to this panel because we have extensive methodologies that have been used both within the financial services sector, within the financial market sector, and also through the many other sectors that IBM serves.

So, again, thank you, and I appreciate the opportunity to participate.

0033

MR. ROSS: Hello. My name is Jon Ross. I work for -- am CTO for GETCO, LLC. I've been there for about five years. Previous to that I was the CTO for the NASDAQ Stock Market, and I was the executive responsible for their single book integration of SuperMontage INET and BRUT platform. So I've got a few big software rods.

My background is software development. I worked at Microsoft, did some videogames in the past.

GETCO is a market maker. We primarily make two-sided markets in the public market. That's our primary strategy. We also run a customer service business and our footprint in the market is fairly large.

You know, our view is the markets are going a sort of multi-decade transformation from a largely manual market to a largely automated market. They're also going through a transformation in complexity, but this isn't happening in a vacuum. The SEC has periodically visited the growing technology of the market, and most notably the implementation of ARP in '89. That was clearly wise given the events and what happened after that; expanding ARP to ATSes in '98. So I think it's highly appropriate that the SEC is visiting this once again and hopefully will shepherd us into a more stable system.

I want to echo some things people said. As our markets get more complex, we're also driving a great deal

0034

of value out of them, but as the complexity goes up, there's risk involved. Anyone who drives a car or uses a computer understands that cars get more complex. Computers, the Internet, all gets more complex. They tend to be prone to errors and disruptions.

Minimizing that is critical, and then as the Commission realizes that, and also when it does happen, dealing with it in the most appropriate manner and dealing with the built in conflicts in management crisis I think is a very important topic as well.

Thank you.

MR. R. COOK: Thank you.

And, again, thanks to all of you for agreeing to be here today.

COMMISSIONER GALLAGHER: Hey, Robert, can I interrupt you for one second?

MR. R. COOK: Yeah.

COMMISSIONER GALLAGHER: In your introductions earlier you mentioned the Commission. I was on the phone before I came down with Commissioner Aguilar, who has been very ill and barely has a voice, but wanted me to tell everybody that he's watching on the Internet the Webcast. So big brother is watching, and my guess is that he'll send us an E-mail if he has any questions, but I just wanted to point that out to everybody.

0035

MR. R. COOK: Thanks. I appreciate that, Commissioner Gallagher, and sorry to hear that he's not feeling well today.

So again, we appreciate your willingness to spend the time with us today and to share your perspectives on these issues.

I just want to note for the audience here and those who are watching that most and perhaps all of our panelists have actually submitted written comments to our public comment file with their observations and recommendations to the Commission, and these are available for anyone to review on the Commission's Website, and I want to take this opportunity to note that the comment file for this roundtable will remain open, and we invite any interested parties to add their thoughts and observations which you can do through the Website.

We have already received a number of very thoughtful comments and recommendations in this way from a variety of different perspectives, and our staff will be reviewing all the comments we receive in the file.

So with that, let me ask Dr. Leveson if you would like to begin with your sort of stage setting remarks, and we'll have to pass the mic down.

DR. LEVESON: Thank you.

0036

First of all, thank you for inviting me here today to share my experience. I apologize in advance that I'm going to have to disappear early. I have sort of triply booked today.

First, let me give some additional background that wasn't clear from the official bio, and that's that I've been in software engineering now for 47 years. I

first worked as a system engineer for IBM, and then I went back to school for a Ph.D. in computer science, and since then I've been teaching and doing research in software engineering.

I've also been a co-owner of a software company for 20 years. So I have a lot of experience, and I've worked basically in every industry in the world. I was originally in the computer science department, a professor for the first 18 years of my career, and then moved to aerospace engineering for the last 14 for a variety of reasons that are irrelevant today, but I am not just an aerospace engineer. In fact some aerospace engineers say I'm just a computer scientist in disguise.

So let me tell you a little bit of what I've learned in the last 47 years. The first lesson is that all software contains errors. I have not in all of that time ever come across any software that was not trivial in which no errors were found during operations. The

0037

errors may not surface for a long time, but they're lurking there and waiting until just the right conditions occur.

There are also some myths about certain industries being able to create perfect software but unfortunately this is patently untrue. No industry creates perfect software. So let me tell you some examples of things that we've -- you may have heard.

So, for example, the space shuttle was rumored to have zero defect software. I always chuckle when it's always called zero defect software. When I chaired a committee for NASA, they examined the software processes for the shuttle software in 1992. We had already by that time found 27 safety critical errors in the shuttle software that could have brought the shuttle down under certain different circumstances, and they were sort of lucky. None of these, of course, led to a disaster, but at that time, just so this may be of interest, 1992, NASA was spending $100 million a year to maintain that software.

Now, most places, even government agencies can't spend $100 million a year just on maintenance of software, and that was relatively simple software. It was only about 200,000 lines of code, whereas jet planes now, commercial planes have about five million. The

0038

military aircraft have 15 to 20 million. Cars have upwards of over 50 million lines of code in them today.

So we're talking about simple software, and they were spending as I say to get such good results from their software $100 million a year; probably spent more in later years.

What about newer spacecraft? That was an old one. Well, as the use of software has grown in space so have the problems. NASA pours an enormous amount of

money into software and employs the best engineers in the world. They know that their projects take upwards of a decade, ten years to create, and they get one chance and that's it. If they fail, it's all over.

For example, every -- oh, and with all this effort and expense, they still have lots of errors and losses because of software. For example, every Mars mission has had serious software problems, and some of them occurred when the software could be put to sleep. What they do is essentially put it to sleep until the NASA JPL engineers on the ground can figure out how to fix it and what to do about it, and then they send up a patch to the software.

In other cases, for example, the Mars Polar Lander, the Mars Climate orbiter and others, it wasn't in a situation where they could put it to sleep and they

0039

just lost everything.

So what about aircraft? All right. We all fly on aircraft. I don't mean to make you scared. We do a pretty good job with aircraft, as you know. We have a wonderful safety record, and most accidents that do occur are blamed on pilots because it's easy to blame pilots. There's a lot of reasons why, but the fact is that most of these accidents blamed on pilots at the very least, the software helped induce the pilot error, and sometimes the software more directly created the problems.

We're probably been reading about the recent Air France 447 loss when it was flying from Rio to Paris, and that is, again, blamed on the pilots, pilot training. As it turned out it was really a technology and a software problem down underneath.

And remember there are tremendous efforts to make this software failsafe and fault tolerant. Air traffic control centers have had outages for hours and had to shut down operations at airports and just tell the pilots to go wherever they can and land because they have unexpected problems.

I could go on and on and with other industries. The stories are all the same, and all of these losses, the software was constructed by highly intelligent, highly skilled people who thought that they had adequate

0040

backups and that they had prepared for every contingency. The very highest software engineering and technology standards were used. They did everything that we know how to do or tried to do.

And things are unfortunately getting worse, not better, and why are they getting worse? Well, because we keep wanting, the way humans are want to do, we keep wanting to make things more complex, get more functionality, do more things with computers. They're almost considered magic devices.

And one of the things about computers and on

software is we have pure design, design that's abstracted away from the physical realization of the device. So we can build things that are much more complex because we don't have all that physical constraints that we had to deal with, and so what we're finding is that we can't anticipate all the potential unsafe and dysfunctional interactions among the components. It's not necessarily just individual component failure. In a lot of these accidents each individual component worked exactly the way it was expected to work. It surprised everyone in the interactions among the components.

At the same time our attempts to create failsafe or fault tolerant systems that already use human monitors or back-ups haven't been terribly successful for

0041

a variety of technical and psychological reasons. I don't have time to go into them, but basically it's almost impossible for humans to monitor computers.

One of the reasons is that they're too reliable. If they failed more often, it would be much easier to monitor them, but if something only makes a mistake every six months or every year, how do you keep alert? How do you maintain the ability to catch errors?

And after a while there's this thing called the incredulity response. We just don't believe that the computer, since it always does the right thing, that it could possibly be doing something wrong, and so people are loath to intervene. I mean, I'm not sure what happened in Knight trading, and it would be fascinating to talk to the people, but I have a feeling that some of these well known principals were probably involved there.

Kill switches, panic buttons, we put those in all of our systems. Unfortunately they don't always work, and sometimes, in fact, our protection features are the actual things that bring us down. The problems are in our attempts to provide extra protection.

And I haven't even mentioned the problems of security and software, which you're all aware of because we all deal with those all the time. So I'm not even going to talk about that, but you know it's almost

0042

impossible to build. Almost? It is impossible to build totally secure software systems. The Pentagon has been trying for a long time, and they have the best minds on the planet working, and there are still people breaking into their systems. They don't like to talk about it, but it's happening.

But some groups have been more successful than others. So we can learn from the groups that have been successful and that have tried very hard to do things right. Of course they use the best software engineering principles. They clearly think of quality, use high quality assurance methods, testing, others. I mean, they just do the ultimate of what we can possibly do, and they

take sometimes a decade doing this before, for example, in the spacecraft software.

So while I'm not suggesting that anyone shouldn't use the highest standards, it's not going to be enough. I wish it were. It's not. So how do the most successful industries limit risk? They use three practices in addition to the standard technological practices.

One of these is oversight. Most industries that have very high reliability software have government agencies overseeing what's being done with an iron hand. For example, the FAA and the NRC, Nuclear Regulatory

0043

Commission, the aircraft industry knows that people will stop flying and we finally would have high speed rail in this country if the planes start falling out of the sky, and if they don't exercise the utmost discipline in creating and maintaining aircraft and air traffic control software.

So they participate. They understand that public confidence is a critical part of their industry, and they do as much as the FAA. The FAA is really working with them in partnership. It's surprising, and the industry itself writes its standards and does a very good job of it because they want to stay in business, frankly.

The nuclear power community until recently just didn't allow software in the nuclear power plants. Unfortunately, we're now starting to build fully digital nuclear power plant protection systems, and the Nuclear Regulatory Commission is starting to provide strict standards and oversight. One of the meetings I'm supposed to be at right now -- I sent a grad student instead -- is with the Nuclear Regulatory Commission to help them figure out what they should be doing. It's one of my best graduate students though.

MR. R. COOK: Yeah, feel free to leave.

DR. LEVESON: I promise.

0044

(Laughter.)

DR. LEVESON: One of the wonderful parts about being at MIT is your grad students are all brighter than you are.

Other agencies haven't been so on top of things, and they all are starting to be concerned, such as the FDA, which is the other place I'm supposed to be right now. You would be appalled if you knew how many people are killed by software and medical devices every year. I mean no one wants to talk about it, to be honest, because the numbers are large.

The first best practice is that the government oversight has been used in these industries. First of all, they understand, as in your industry, that public confidence is critical for the survival of the industry,

and the government and the industry have worked hand in hand to ensure that public confidence is not disturbed.

The second practice they use is essentially being extremely conservative in the technological devices. Now, it's not that they don't use the latest technology. They do. They use the latest and greatest technology, but they limit software functionality and complexity. So the software in these successful industries contains for the most part only the minimum function that's required to achieve the goals of the

0045

system. It doesn't add on stuff. You don't mix up the entertainment software on your airplane with the other software.

And unlike the financial industry with its high frequency trading, the other industries that require public confidence to survive limit the complexity of the software they build and they only implement, as I say, the basic functionality they need to get away with to serve the primary mission. Any extra stuff, entertainment systems and stuff is strictly separated.

People want to do that; you want to have entertainment on planes, absolutely important, but you don't mix that up with your avionics software.

The third and final practice I want to talk about is the application of systems thinking and system engineering. These industries realize the problem is not just a technology problem; that they need to design the larger system so that software errors don't cause mayhem because they know that the software errors are going to occur despite what they do.

And they do this through providing the control structure that limits and controls risk by enforcing constraints on the non-technology related system behavior that preclude or at least greatly reduce serious losses. They understand that they need to fix the system, not

0046

just fix the technology.

I wrote about how to do this in my new book. I'll get a little plug: Engineering: A Safer World, which was published last January and is getting a lot of attention, and it's really a different approach, but it talks about how do you apply systems thinking. How do you look at the larger system and fix that, too, along with the technology?

So sort of to summarize, I don't want to sound like Chicken Little or a latter day Luddite. I did not get a Ph.D. in computer science and spend 47 years working in the field just to try and convince everyone not to use computers. But the bottom line is that there's 100 percent certainty that you will have more upsets caused by the financial system software and probably in not too long a time, but it will occur, and it's probably going to start occurring unless something

is done more frequently because people are going to keep adding more functionality and more risk into the system unless it's limited.

And there is no technical fix. It doesn't mean we shouldn't test, as people were talking about. Use the highest quality assurance methods, do everything we can to build great software, but that's not going to totally solve the problem.

0047

The industries that have learned this lesson the hard way limit their risk with discipline and establishing controls. The biggest mistake that the Titanic designers made was believing that they could build an unsinkable ship and, therefore, they didn't have to prepare for contingencies, for calamity.

And this is still true. We've learned better in these very safety critical industries. We've learned that we cannot build an unsinkable ship and we cannot build unfailable software, perfect software. The financial industry needs to learn, too, that computers aren't magic; that our engineering techniques for creating software aren't perfect; and that failsafe and fault tolerant designs, whether these features are automated or they use humans in a monitoring function, are a goal but not yet a reality.

We need broad approaches and solutions that go beyond the technology. If instead this industry engages in hubris and wishful thinking, we're all going to have to live with the consequences.

Thank you.

MR. R. COOK: Thank you, Dr. Leveson. We really appreciate your insights, and I think they provide a valuable backdrop and context for the rest of our discussion.

0048

I know you need to leave. So we appreciate your spending your time here. I suppose there are probably a number of questions those of us who aren't flying today might like to ask, but in the interest of making sure that we have time to hear from our other panelists, I'm going to move us on to that unless there's any objection, if anyone has any.

Again, thank you for joining us today.

DR. LEVESON: Thank you.

PREVENTING ERRORS THROUGH ROBUST SYSTEM DESIGN, DEPLOYMENT, AND OPERATION

MR. R. COOK: In some respects I guess, you know, your comments cast new light on the title of this panel, "Preventing Errors," but we're eternally optimistic, and so we will begin to explore this through a series of questions that we have for our panelists.

And the first one is really about industry best practices, and maybe I'll ask Christopher to take the first crack at addressing this, given your kind of

perspective working with a variety of different financial services firms.

But can you talk to us a little bit about what are the industry best practices for testing robustness, deployment, and the use of software systems? In particular, are they sufficient today to support the

0049

continuity and integrity of the markets? And if not, what more from your perspective needs to be done?

MR. RIGG: Sure. Well, first I would say that I think when you talk about best practices, I think best practices are aligned to particular industries, and different industries have different levels of risk that they face in the marketplace, and that tends to drive the practices that they choose.

So just as Dr. Leveson talked about when the risk of what you're deploying into production goes up, the amount of rigor typically that you apply to that process should go up commensurate with that. So higher risk applications should have more process rigor associated with them.

So I think from a best practice perspective we're seeing that through many industries where you can see in the financial services industry you have very large, complicated systems that have to be up all the time. If you look at the Fed. wire system, if you at ATM systems at large banks, those systems have a great deal of rigor associated with them, which means very small number of changes, an extensive amount of testing prior to those changes going into production, and a significant amount of sign off by the various stakeholders.

So I think, you know, it's hard to apply a

0050

single practice to any particular industry because obviously some industries have a great need to have a high frequency of change because they feel like they need that in the marketplace in order to respond to the competitive nature, and I think this industry is certainly one of them where there's a significant amount of change. So the tendency can be to use less rigor when there's more frequency of change.

I would say in the best practices space you need to continue to figure out a way to balance those two, the amount of rigor with the frequency of the need to make change very quickly.

In terms of whether they are adequate, I think, you know, I'm not sure I'm in a position to determine that. I think that every business needs to look at their processes, the key stakeholders. One of the best practices, I think that we all who develop software know that you need to make sure that the key stakeholders, which typically includes the business, whoever owns the P&L, they need to be integrally involved in the process so they understand what they're asking for, the risk

associated with what they're asking for, sign off and agree to the amount of testing that's going to be applied because ultimately at the end of the day all of us here who are technology people, we typically don't run our businesses. We are providing the technology that the people who own the businesses are providing.

So in that case I think one of the key things there is to make sure that the business understands the risk associated with the technology being developed and they agree to the level of rigor that's being applied.

CHAIRMAN SCHAPIRO: Robert, could I jump in with a question?

I thought that was a great presentation from Dr. Leveson, and one of the things that really struck me was this idea that in the Nuclear Regulatory Commission world or the FAA, they're very conservative about their use of technology, and they go towards, you know, just dealing with the essential functionality and not lots of add-ons.

And I wondered if in your experience at IBM and really for anybody in the financial services industry has that discipline of let's focus on exactly what must happen with this software and not throw in lots of bells and whistles, which as she describes it leads to the potential for the interaction of components that creates problems.

MR. RIGG: I'll start and encourage other people to contribute.

I would say that I think in most cases for key mission critical applications that have certain requirements associated with them, whether it be high speed of execution, obviously which is something that's common here, I think in those cases I think most of the time the software is minimized to what it actually needs to be, and often that is simply because of the requirement to execute at the speed that they need to in these marketplaces.

So I think if you were to look at the core routing algorithms and core processing software of any of the people up here and any of the participants in the industry you would find that that software is fairly lightweight, fairly focused specifically on the task and doesn't have a lot of extraneous feature and functionality.

I think there are other parts of, you know, the organization of many businesses that have software that's more feature rich and more focused on that because they feel like they need that in the marketplace, but I think for the key areas of financial systems that are performing very specific tasks, I think that that minimization principal is generally applied, and what I've seen is it applied pretty well and pretty

appropriately.
MR. ROSS: I have a quick comment. The most
0053
stable platform I ever worked on, which was also the simplest, and I don't think that's coincidence, which was Ireland, just to bids and asks, matched them. No opening crosses, no pegged orders, nothing fancy at all. This was obviously pre- Reg. NMS. That was by far the most stable platform I ever worked on.
That platform we decided is not appropriate for today, you know, with all of the interlinkages between exchanges, open and closing processes that need to be implemented, but I think it's maybe worth -- I think it's maybe worth reflecting on for a minute how much of the complexity in the system do we actually need, to your point exactly.
MR. LAUER: I'd like to say that I think sometimes from what I've seen we can go too far actually, especially in low latency systems, when we focus simply on what the task is for that software to perform. I've seen risk checking being neglected in favor of faster performance, that type of thing.
So on one hand, you can go too far in simplifying too much, especially when there is this latency race to zero. On the other hand, it can be very tempting by designers of these systems who are generally very smart to see them as much simpler systems than they actually are, and so when they go to make changes to what
0054
they perceive as a simple system, they don't always understand the complexity of even minor and small changes, and especially the complexity of how these algorithms are interacting with one another in the marketplace, and that can lead to much less rigor in terms of testing systems both from a quality assurance perspective where you tend not to have independent quality assurance groups in many firms, as well as a load testing perspective where software is not adequately exercised in conditions that, you know, accurately mimic what the market conditions are like when you're actually trading.
MR. ARYA: I'd like to add to your point. I think isolating smaller components is extremely important, but I think also taking, to kind of complement that, taking the approach to have a multi-shell approach so that you will have smaller components that are written with simplicity, but in order for risk checks and so forth, we take a very multi-layered approach.
So you have other risk systems they are continuously observing how these low latent systems are trading, and they have means to intervene in real time. And it's extremely important to keep these low latency systems very, very simple, but also it's very important to have integration between these outer shells that

0055
protect these simpler systems and continuous integration testing between the two.

We run a dark pool, one of the largest dark pools, and we specifically go out of our way to make sure that the actual crossing engine itself is extremely simple, but that said, its interaction with tape recording, its interaction with order writing and so forth has to be continuously tested. The crossing engine could be all doing well, but if the external components are not doing well and interacting properly, that could lead to a disaster.

So having smaller pieces as simple as possible but having outer shells that protect them, watch them is equally important.

MR. NAZARALI: If I could just make a point, we're spending a lot of time talking about software development and testing and implementation, and that's very important and all of us can do a better job doing that. But as we think about trying to avoid the type of massive errors that happened at Knight Capital Group, it's important to think about the whole system.

And as Dr. Leveson pointed out, two things really resonated. Number one is all software contains errors, and number two, it's not just a technology problem.

0056

So if you look at what happened at Knight Capital on August 1, the first five minutes of that trading was really a software problem. The next 35 minutes where the software was not shut off was really a risk management and control and management processes problem. And I think it's important for all of us as we make our systems more robust and improve how we implement the software, that we also put in place the right management and risk protocols so that if something like that happens we can pick up the phone and call the New York Stock Exchange and say, "Shut off all trading. Kill all open orders."

If that had been done five minutes after on August 1 at Knight Capital, we probably wouldn't be here right now. So I think that that's something that's very important for us to all consider.

MR. PASTINA: Just another point that Dr. Leveson mentioned was that the amount or number of lines of code in a car, and that really struck me because she was talking about millions of lines of code in a car, and I don't know about the other participants, but a typical matching engine doesn't have millions of lines of code. It typically has thousands of lines of code. It's much smaller. It's much more streamlined. It's not anything like that, and you know, comparing the two is very

0057

interesting because the amount of testing that goes into

a car compared to the amount of testing that goes into a matching engine, it's a much smaller component. It's not as big as, you know, a couple million lines of code.

MR. ARYA: I think in terms of best practices, I would like to comment that it's really -- Dr. Leveson said that errors are inevitable. I fully agree with that, but I think accepting that and understanding that errors are inevitable and actually really cultivating a prevention culture in your team is extremely important.

Before the first line of code is ever written, what design principles go into play? To your point, we think that if technologists are really interacting with the business folks, folks that manage P&L, and they say, "Okay. I'm writing a line of code where I'm about to see an erroneous trader, an erroneous code. What do I do?" and if they're not integrated well with business, their temptation or instinct might be, "Well, I'll throw an alert and good things will happen."

But if the two teams are very well integrated, they will go and talk to the product manager. They will go and talk to the business guys and say, "What's next?" and that really should go into our design phase, talking about what error conditions happen.

Everybody talked about errors, and they will

0058

happen. What really gets us, to Jamil's point, is double errors. Software errors happen, and there's a cascading effect, and that really needs to be integrated into our software design.

Error prevention at first level is very possible, mostly doable, but it's the double failures that are not planned for, and we try to induce that culture where you really talk about there was an exchange issue or a broker issue on the other side. Why do some of the networking errors happen? Or how do you prevent?

You try to pull off a kill switch that is internal for a client, and it didn't work. So double failures really need to be built into these best practices and need to be thought of when you design the system before you write the code ever.

MR. R. COOK: I think that's an important point and Jamil made the point earlier about it's not just the front end. The back end is sort of in a way contesting the logic of the two panels we have because they do overlap in certain respects, and I think we recognize that.

But coming back to the best practices question for a minute, maybe you could give us a little inside window into is that even the right term in this industry. Are there commonly understood best practices?

0059

A lot of you have worked in different organizations. As you move from one to another, does the best practice of that organization for developing code,

for testing it, for rolling it out, does it look similar to where you were before?

When you talk horizontally across different types of market participants, is the idea of best practice sort of -- does it exist regardless of the nature of the entity you're working for? Does it exist at all?

And I'll open that up to anyone on the panel.

MR. PASTINA: If I can jump in just quickly --

MR. R. COOK: Sure.

MR. PASTINA: -- I mean there is a system development life cycle. There's an idea. That idea has to be documented typically in terms of specifications. Those specifications get turned into, from a business set of specifications, into a set of specifications that a coder can understand, so a design set of specifications.

There's a coding phase, and after the coding phase there's a testing phase and implementation phase. So that cycle exists everywhere. However, how that cycle is implemented is widely different. So in a more entrepreneurial type environment, the specification is written on a napkin, right? Someone has an idea in a bar

0060

somewhere and they say, "This is a great idea. Here's the idea. Here's the spec coded."

And it goes right into coding. Somebody looks at it and says, "Yeah, this is about what I expected to see coming out of coding," versus a more disciplined approach where an idea has to pass through a set of filters. It then becomes a specification, and the least expensive place to weed out errors is early on in the program. So the most expensive is obviously in production, testing, coding, but the least expensive is in the specification phase.

And so there are techniques for inspecting specifications in a very disciplined way to weed out errors up front where it costs you very little to do.

MR. NAZARALI: If I could just make a point, we do believe that there are software best practices, and they really fall into a couple different categories. One is software best practices just in the development implementation of software, so things like, you know, unit testing and regression tests and things like that.

The second set of software best practices really relate to very specific firms. So, for example, our software development practices in an automated market making business are going to be very different than the software development best practices for an ATS or for an

0061

agency trader.

One of the really powerful things that came out of what happened on August 1st is that Wall Street firms across the industry said to their technology teams, "We want to make sure this never happens to us. So go back

and, you know, make sure that the system integrity and how you roll things out and the checks are such that the probability of this happening is miniscule."

And as an industry, we met with other firms. We met with both competitors and customers, and we shared some of the best practices, and many of them were things that we were already doing and things that they were already doing, but there were some ideas there that both we and some of the firms that we met with found very, very helpful.

For example, at Citadel, we have developed this fuse box technology over the last ten years, and it's a system that sits outside of your trading software, and it listens to all the trades and executions that happen, and under certain conditions much like your fuse at home will trip, if the ADV that you're trading is much higher than a set parameter, if your risk limits are much higher, if your P&L goes out of bounds, it will trip and then a human will then have to turn it back on.

So things like that were not necessarily things

0062

that other firms were using, and that was very helpful as we shared best practices.

MR. ISAACSON: I'd like to also mention as far as best practices and as I mentioned in my opening remarks, I think everyone here is probably doing unit testing of the different components, but something that Dr. Leveson mentioned was that usually where the errors arise are from the interactions between components that are operating perfectly. So integration testing and regression testing is paramount.

And you can do that within a firm, and it's probably very similar within firms whether you're an ATS or an exchange or a market making firm. However, what's very difficult is to simulate the market. How does the market actually interact with 13 exchanges and many dark pools?

And that's why I'm advocating that all market centers support test symbols, and then each strategy is rolled out using a test symbol in production environments where firms like Jamil's are interacting. You know, they're testing their new strategy in real live production environments where real live quotes are being taken, but using a test security that doesn't clear, that's not printed on the tape, but everyone can see all of the interactions, especially, you know, the designer

0063

of the system.

And the exchanges can do this as well with new order types. That's really the truest test. That's the final test once you've gone through unit tests, integration testing, and then user acceptance testing. This would be the final test where you're testing with test symbols in a production environment.

Now, realizing that errors will occur and we need to have things on the back end to make sure that they're, you know, remediated immediately and stopped, but I think that would go a long ways in our environment.

COMMISSIONER GALLAGHER: Can I jump in on that real fast?

Chris, I guess the question is you just can never replicate the real market for testing purposes. You can take, you know, a reel from yesterday's trading and test into that and expect that it's as good or vibrant as --

MR. ISAACSON: I'm not going to say never. I'm saying it's an incredibly difficult thing to take in all of the market data feeds and in perfect sequence regenerate them back and consider how they would interact with your matching engine or your algorithm.

But you know, that's why you need test symbols to test that new strategy in the real live environment.

0064

MR. LAUER: I'd like to say that that doesn't mean that you don't try to mimic what a real market looks like, and in my written comments I talk a lot about this, which is the difference between software replay and hardware, which can do what's called temporarily accurate replay.

Market data comes in in the modern market very fast and is subject to conditions that are called microbursting, and that is when multiple servers can write to the line at once resulting in temporary network saturation, and that can lead, when you have complex multi-threaded software, that can sometimes lead to nondeterministic behavior with thread contention.

A very technical conversation, but the main point is that many firms do not use hardware replay for temporarily accurate replay, and the ones that do, maybe they do so initially when the strategy is designed, but not as it is changed over time. For most firms -- and this is especially according to the Chicago Fed. report -- you know, they found out that a lot of firms will make small changes and constantly push stuff out, or if they come up with a new strategy, they'll back-test it and say, "Hey, this is going to make some money," and put it into production very quickly.

In terms of a software development life cycle,

0065

I do believe there are best practices that can be adhered to, and exactly as was being described, its documentation specification is the step that is normally not taken at least for many firms.

There are certainly mature firms -- Citadel and GETCO are great examples -- of very mature firms with very mature processes, but we live in a world in which a single server can send out 100,000 orders per second, and so we live in a world in which small start-ups and small

firms that want to gain an edge and want to move fast will definitely lean towards not adhering to a software development life cycle and not properly testing and dealing with the types of things that you asked about in the preparation, which was the behavior and the unexpected condition testing.

And I think that there are things that can be done in terms of identifying certain times, certain market events that should be mandatory testing perhaps or saying that if you do have low latency systems, they do need to be testing with hardware and temporarily accurate replay data, which is a major step.

And then I completely agree that life testing in a fragmented market with test symbol ranges is a fantastic idea as well.

MR. R. COOK: I think Commissioner Paredes had

0066

a question, and then we can continue.

COMMISSIONER PAREDES: Yes, and it really feeds into the discussion so that it may inform what your answers are going to be, those of you chiming in, and that keeps us on the theme of best practices. And a lot of general categories of issues have been identified, documentation, testing and all the rest, but as we try to get more granular, how do we know whether or not something is a best practice or not?

And if you think about testing, how do you know whether or not there's been enough testing? I'm sure you can always anticipate. Well, we can do one more round of testing, one more round of testing. We can always do a little bit more.

So in part it strikes me it comes back to questions of risk tolerances, as well, and tradeoffs and the like, but identifying and finding agreement that there should be best practices or there should be testing and there should be testing of this type versus that type, how do you all go about evaluating it?

As a business matter, it strikes me it's important for policy makers, too, to have the benefit of that kind of more granular insight to figure out, all right, at some point a decision needs to be made as a business matter, as a regulatory matter that to do yet

0067

another test, to implement yet another practice, to go yet another step isn't sensible given the set of tradeoffs.

So from a business perspective, how do you know; how do you think about enough is enough; now it's time to go ahead and roll it out?

MR. NAZARALI: I think you never really know if the testing is enough, and I think it's really just a business judgment where you look at the cost benefit of the potential error. So, for example, Citadel Execution Services is a large retail equity market maker and we

roll out software changes all the time partly in response to regulatory requirements, partly in response to customer requests, partly in response to changes in our trading strategy, and whenever we do a rollout, there is a decision made where: okay, should we wait; should we do more, as you said, more testing, or do we implement this?

And it's a cost benefit where you say: okay. There is a one percent chance of this creating a problem. How big a problem will it create? How many customers will it impact? And, you know, what's the magnitude of that? Okay. That's unacceptable.

And you never know that percentage exactly because you may think it's one percent, but it could be

0068

three percent or half a percent, but it's really that kind of cost benefit where you try to anticipate the potential cost of something going wrong.

CHAIRMAN SCHAPIRO: And how do you think about the cost not just to your customers, but to the broader marketplace, investor confidence, and the customers of other firms that may be impacted by your judgment that we're ready to go?

MR. NAZARALI: that's a great question, and for us because we are such a large player in the market, we do consider all of those things on the market making side. You know, we have roughly 25 percent market share in retail equity markets. So we realize that if there is an impairment of investor confidence, it's going to affect our business.

The concern though is that smaller firms that don't have such a vested interest in insuring the integrity of the markets, they will make a very different calculus because when they're rolling something out, the cost of implementing investor confidence and hurting the overall market doesn't affect them nearly as much as it would affect someone like Citadel.

MR. ARYA: I would like to add that regarding testing, as Jamil said, there's always a line you draw, and it's judgment based, what releases go out, and it

0069

really needs to be, as he said, you know, talked about between technology and business folks.

But within the testing itself, it's not just the testing of the software itself, but it's also the testing of how the software will recover if there are problems.

So you may never ascertain that there are no errors, as Dr. Leveson said. However, it's not just testing that the software is functioning properly. We also look --when I say we're okay to go ahead with this release, we also look at what testing has been done about failures, what testing has been done with respect to diverting the software back to the previous version.

So it's not about that the given version would work or not, but it's also about if there are unforeseen issues, how quickly can you divert back.

To Chairman Schapiro's point, investor confidence and so forth, there will always be error, and we can make an error of judgment, but as long as you code and you test for diverting and curbing those problems as quickly as possible, and that needs to be tested, and your code switches need to be tested and your intervention needs to be tested, and if those are tested, then you have some degree of confidence that, yes, you can divert back. Errors will happen in these extreme

0070

cases, and there are ways to interact with other systems and how you recover and how you quickly go back.

So deployment of that software and diverting of that software is equally important from a best practices perspective, as is the testing, because there's the next phase.

One last thing I want to mention. It was mentioned about exchange testing as well as back-testing. One of the areas where we have really got some traction is in operating our own dark pool and crossing engine. We had previously struggled quite a bit about rolling out a new software release for the crossing engine and having issues.

So over the years we have actually built a parallel production testing environment where all the real time flow is also directed to the next version of software. It trades in parallel. It compares in parallel. It is actually subjected to the same kind of market data, and I think more and more firms, even in smaller scales and not everybody has the luxury to be a dark pool because it's an endpoint, so I totally recognize that for exchanges that problem is a lot harder.

However, it's really paramount that, you know, some of us get together and devise those ways or talk

0071

about how this parallel testing can be done because really the only way to subject your new software to real time market conditions is to actually subject it before you release it, and it's an interesting computer science problem, but that can also be worked out with talking to industry guys and business folks.

MR. R. COOK: So I think we have questions from Craig and Tom, but I'd like to just wrap up the answers to Commissioner Paredes and Chairman Schapiro's questions.

MR. ISAACSON: I just have one more thing about that all important decision where you're making a judgment call. Is the software ready to be released? It's not just the business involved with technology, but it's also Compliance Department, Regulatory Department.

There's an exchange. Does this new functionality match what our rules say?

And so you've got to have not just business and technology but compliance involved in order to really make the right risk-reward tradeoff.

MR. LAUER: I'd also like to say I think that Chairman Schapiro's question is the crux of the matter. Does an individual firm adequately account for the impact it can have on the market and the broader implications?

Which is why I believe there's a clear

0072

regulatory role as well for any firm that has direct market access. That's the enter point where you can start mandating potential quality accreditation standards or allowing firms to use quality accreditation in their best ex decisions so that there can be at least in that regard a private market sort of race to develop better technology and better quality management standards.

MR. ROSS: We've been talking about testing of individual components and when is it time to release a component. If I have a component that I'm changing, the component in front of it that's doing extra-regulatory checks I don't change, right? So I have a chain of components all doing checks. I only change one at a time.

If I'm changing the one in front, I don't change the one in back that I know behaves properly, right? So nobody should rely on a single component having checks.

To that point, one thing that no one has talked about except for Mr. Ratterman in front is the financial industry is unique in that you can actually see your effect on the market independently. So we have real time drop copies for most of the exchanges. What we do as a firm is we have separate code that implements those drop copies and separate code that reconciles what we think we

0073

do in the market from our proprietary trading systems to what we're actually doing in the market.

Almost no other industry has this. You can't have a robot that flies alongside a plane and when you push the stick down sees the surfaces turn. This is very unique. I think it's very important.

So from a best practice point of view, I think there's a very key best practice. It's not necessarily in software development. It's in how your system operates and how your system reconciles. To be slightly pejorative, some exchanges make money off of their drop copies, which might not be the best thing. Also some exchanges drop copies are actually not real time. They're like way back in system. So your view from the downstream drop copy is just slow enough that you can't trust it. I think those would be interesting areas for the Commission to explore.

But that independent view is critical, and I think any institution worth its salt spends the resources to build that independent view of what effect they're having in the market in real time.

COMMISSIONER GALLAGHER: Before we leave testing, what are we doing now that impedes proper testing or, on the flip side, what could we do to facilitate a more vibrant testing atmosphere?

0074

MR. ARYA: I think the point was made about live symbols. That would greatly help us as an agency broker. It would greatly help us, and I think the key point is that all exchanges should participate in that, and also we should encourage market participants because exchanges could provide functionality to test your symbols, but it's really a function of how much feeds actually go into it. So that's extremely important.

And if there are ways to -- also I think the point about drop copy is paramount. We think that there's a lot that can be in the drop copy and done in the drop copy area; would love to find ways to make them real time for all exchanges; would also love to actually find ways to get aggregated alerts of some sort where I think one of the comments had that exchanges could have peak notional values and so forth and threshold goals and a communication comes back to the brokerage firm.

But what would be very good is if that can actually also be somewhat automated. So if my firm is sending above and beyond 75 percent of what I typically send, whether it's ADV or peak notional value, having a feed in drop copy warning these firms in real time would be very, very good.

And, again, guidelines to do that and industry working together would be extremely helpful because we

0075

can actually test that and also have our own alerts in sync with what the exchanges are seeing.

MR. ISAACSON: I think longer term there's another thing the SEC can do, and that's the forcing of real time reporting to the DTCC. We at BATS believe that while there may be interim solutions at the exchanges where we have kill switches based upon limits at each individual exchange, ultimately that functionality belongs to the DTCC, and in order for them to do those real time position limit monitoring, they need to have all the trades in real time. And I believe they have the vast majority of them in real time today, but not 100 percent.

And like drop copies from the exchanges, a firm can't get a full view, an independent view, without 100 percent drop copies, and likewise DTCC can't make a 100 percent calculation of the exposure position of that member.

COMMISSIONER GALLAGHER: I'm not going to

disagree with you, but I'd point out that I think their pricing structure has to facilitate real time.

MR. ISAACSON: I would agree with you. I would agree with that. We have no interest in, you know, necessarily increasing the cost of clearing materially, but I think, you know, this is an important enough risk

0076

management discussion that we should make both changes to make it happen.

MR. NAZARALI: This may bleed into this afternoon's panel, but since you're asking what you could do, one of the things that we think is really critical and, as being part of the industry working group that submitted the recommendation, is we think that kill switches at the exchange are very, very important because as we've talked about here, we can all improve our processes, but software will always have errors, and our trading systems may -- things may slip through.

And so the exchanges as the gateway, as the final stop where the trading happens are really in a position to be able to shut off this aberrant trading activity if they see it, and we think that's very, very important so that a problem at an individual firm doesn't become so large that it threatens the stability and integrity of the overall market.

MR. LAUER: I have to say I believe that exchange kill switches are an important first step, but a very inadequate one. I think there is no way to approach the problem of kill switches without a market-wide perspective. Everyone is trading on different markets. No firm sees a market in isolation, and kill switches in isolation can be inadequate.

0077

I think that in my written statement, I submitted an idea for a strategy registration system so that any piece of software that will have direct market access would need to be registered, assigned a unique ID number, and include that ID number in any quote that it sends to the market. While that ID number would not be publicly visible, it would still filter through down to the SEC.

This is not quite the consolidated audit trail. It's what I would think of as an intermediary step in which the SEC could take on the role of market-wide surveillance mechanism, working with exchange surveillance groups. So in a distributed system, but still with a consolidated view, able to keep track of what individual strategies are doing with software, like Dr. Leveson was saying. You can't have humans monitoring software. You need advanced software to monitor software. You need low latency systems to monitor low latency traders.

This is a distinction that is critical, and the one reason that I bring this up in consideration of the

consolidated audit trail is that it is not real time, and while it has lots of features that don't need to be real time, real time surveillance is necessary, and with the strategy registration system, after having back-tested a

0078

system, a firm can tell you what it expects this system to do. What are the characteristics of it?

The SEC can also build software systems that monitor these strategies over time and develop heuristics to actually build what I would call dynamic adaptive kill switches that operate market wide.

MR. ARYA: I would totally agree about the inadequacy of kill switches in a vacuum. So I think metrics for those kill switches need to be really worked out by the industry and finding out there is no one kill switch threshold that fits all. Some firms might want volume related; some firms might want notional value related.

So I think a lot of discussion needs to take place for deciding those metrics. Ultimately I think kill switches should be put in place after that discussion has taken place.

However, regarding strategy registration, I think it might or might not work, and I'm not the expert in the prop side of things; however, for agency brokerage firms which are really writing agency algos, I think a registration and actual auditing of that would be a huge cost to the industry, and I think monitoring of that would be hugely expensive.

Just take into account the consolidated audit

0079

trail. There's a lot to be done even in non-real time putting the data together and so forth. For the amount of work that needs to take place and strategy registration for firms that are continuously evolving and so forth, I think that's a step that will eventually potentially get there, but there's a lot to be done ahead of that, and jumping to just registrations and validations of technologies solely in the game I think is somewhat of an overkill.

MR. R. COOK: I want to make sure no one is prejudiced by not having a mic nearby them. So, Lou, I think you maybe wanted to get in on this.

MR. PASTINA: I did. I had almost forgotten what I was going to say, but going back to the question about what could you do to help in the testing area. It's amazing to me how many times software gets introduced and firms don't test with you. So whether we have test symbols in production or we run industry tests, it's always the same firms that come in and test, and those are the firms that generally don't have issues.

And then there's a long list of firms that never show up. There's also --

MR. R. COOK: What's the balance? I mean, is

it 80 percent show up or --
MR. PASTINA: Well, unfortunately, you know,

the 20 percent that show up represent most of the volume, which is good. However, all you need is someone to come in with a bad message and ruin the whole system, and if you don't have the proper defensive code in place, that message then can be propagated and because of the interconnectedness of our systems today, be passed on to other marketplaces.

And so I'm not advocating that everyone be required to come in and test. I don't think you can wait for every customer to come in before you put a new piece of software in place, but there should be some review about how firms actually take advantage of the testing opportunities that are afforded them.

MR. ROSS: I tend to agree. On the strategy registration, all of the errors we're talking about are not strategies behaving as they're supposed to. The strategies behaving -- or actually they're not even strategies. Most of them are just software behaving as it's not supposed to. It might have unintended side effects on strategies.

It seems like if we could enumerate all of the ways that strategies would fail, then we would just fix them, right? The classic thing: in order to write fool-proofed software, you need a really clever fool, right?

So I'm not sure of the value of registering

strategies because when they behave oddly, they behave oddly in circumstances no one has ever seen or it's a knock-on effect from a straight up software bug. So it seems like minimal returns on that.

MR. R. COOK: Let's get a -- oh, sorry. Did you have another point?

MR. ROSS: Yeah, I just wanted to say, yeah, I think we need to talk a little bit like what classifies as an error, right? Is a firm not being able to connect to a marketplace an error? I mean, that happens every day. That happens to my firm. We upgrade a switch router. We're not able to connect. Since we're not allowed to send any orders, the harm is really de minimis, right? But that was an error, right? We didn't test the switch software appropriately or sometimes the exchange makes a change, et cetera, et cetera. That happens all day every day, right?

So I think as long as those errors are contained to things about not doing bad things, that seems appropriate to me. What we're really trying to focus on is software that is behaving inordinately badly, and my experience has been that is not software that's just rolled out. That's software that's been in production for a while and then hits the magic case that wasn't tested and goes nuts, right?

0082

I mean, that's the case with the IPOs, right? That's even the case with Knight to some extent. That wasn't new software. In my experience every time I have a bug that causes my firm a loss, it's always something that's been out for a while, right? It has gone through all the testing, hit the weird case. Somebody sent a zero price order that reflected off of something else. You know, it's two failures. So I just wanted to make that point.

MR. R. COOK: So, Tom, do you have a question?

MR. BAYER: Yes. So my follow-up question is regarding once you've introduced software and you've got defensive code and other monitoring systems that are available to you, do you go back in and retest your code after it has been deployed? And what are your software testing mechanisms to enable you to do that?

So again, if you could talk about defensive code and how you're modifying your monitoring code according to changes that have occurred inside of the marketplace or your use cases if you're talking technology framework.

MR. PASTINA: If I could just jump in for a second, so quality assurance teams should be imbedded with the specification teams because that's where they build their use cases, and that's where they build their

0083

test cases. And those test cases can reach the thousands of test cases, and for all of us they're all automated today.

And every time a new piece of software gets rolled out, it's not just the new functionality that gets tested, but that software is regression tested against all of those old use cases. So I think that's the way that happens today mostly.

Defensive code you try to build in as best you can. Unfortunately sometimes you don't recognize that you need the defensive code until you've had the situation where you realize now that's where I need a defense. And so sometimes it comes to them late in the game.

MR. ISAACSON: And I would just say similar to what Lou said, you know, we actually at BATS have more than three times as many lines of code that have to do with unit testing as we do actual lines of code in our matching engine. So and that full suite of unit tests never decreases. It only increases. So all previous functionality is tested on every single release we put out in addition to the new functionality.

And we have test driven development that says, okay, this feature doesn't exist. Write a test that proves it doesn't exist. Then write the functionality,

0084

and then prove that it passes the test. I mean, it's

kind of Computer Science 101 or Software Engineering 101.

But you know, we spend an enormous amount of time automating those tests and making sure that the regression never decreases.

MR. BERMAN: Could we talk a little -- could I just ask a question about that level of testing, getting back to something that Jon said?

You're doing unit testing. You're doing regression testing. We all have virus software on our computers. My virus software has never caught an old virus ever because it's always the brand new virus that came out. There's no definition, and then you get the download from wherever and then, of course, it never hits again.

So you build up this massive, massive suite and now my computer is ten times slower because it's checking for viruses that never will occur. So I'm quite sympathetic to the idea that a lot of it is not new software but old software.

Now, there's a parallel in finance that we've seen. In 2006, you never tested whether or not a mortgagee would not pay because all mortgagees paid. There's no such thing as a credit risk in a mortgage, and then a few years later we find out that that was just a

0085

very bad assumption, and we now have massive amounts of stress testing where you have independent personnel saying, "What can go wrong that is independent of the portfolio manager, independent of the traders?"

But really it's really a unit that's almost like a SWAT team. People who are thinking about what to stress, what can possibly go wrong that's really independent of the process, does that exist at your firms? I mean, who do you bring in to say, "I'm actually not part of the development team. I'm not part of the creating the new order type. I sit around all day and I think about this could break if the following things happen"?

If the lights flicker at the New York Stock Exchange at the exact instance that a car rumbles by, this fiber optic is going to pull out and that's going to cause a problem.

MR. ARYA: I think what we thought at our firm, the way we approach it is not by having somebody in a separate room who is basically charged with looking at those error conditions. We tried that to an extent by having a, quote, unquote, separate QA department way back when. What we realized is a separate QA department that's really not living and breathing the operations and understands the code and its interactions with the myriad

0086

processes is actually not capable of really feeding into what kind of issues can happen.

So what we do is we take the monitoring and

administrative staff in the firm that are watching these systems continually operate and administer. They are the ones that are seeing errors happening. They're the ones who are seeing -- as a matter of fact, if you go talk to some of the monitoring staff, they will say when lights go off and others go off, they have a pretty good idea while it's that market participant's that typically goes, does this and this happens.

So folding multiple entities in your firm, product management, administrative staff that sees these errors all day long from other sides of the industry, as well as your QA and desk staff, and coming up with these test plans that are reviewed in connection with all these people giving their feedback is extremely important.

So I think the answer is integrating these teams and giving feedback. That's how we approach it.

MR. NAZARALI: Yeah, we would strongly agree with that. You know, in our experience, it's not very helpful having someone from the outside come into the QA that doesn't really understand the system and how things work.

What's more important is actually having people

0087

that are battle tested, that have seen a number of errors happen before, and when they're putting something in, they remember this hedge condition that caused this other problem and they know to look for that.

So in our experience it's much, much more important to have people that really understand the business, that represent different components, the trading, the risk management, and that have seen a number of problems in the past and are able to reflect upon their experience and draw on their experience to mitigate the probability of that happening.

MR. BERMAN: Now, can those people be somewhat independent, meaning so they have that expertise but they might be from another firm or from a third party, or is it actually you really have to be imbedded in this particular software. You may have seen those errors at another firm, but until you've been battle tested at this new firm with a new piece of software, you're not going to be able to contribute in that same way.

MR. NAZARALI: Yes. You really have to be within that firm because it's so complex that, you know, being from the outside, you're actually going to draw resources away from the guys building it because rather than spend that time testing or implementing it, they are going to spend all their time explaining to you how this

0088

works and going through code by code, and it's really ineffective.

MR. LAUER: Can I just say while it's definitely important to have input from the people who are building software, frankly, if we're talking like

Computer Science 101, Software 101, the most fundamental rule is that quality assurance is an independent function, and one of the gravest mistakes that software developers make is believing that they can adequately test their own software.

So while it is important as always to get input from the software developer and business groups, independent quality assurance, it's a well understood principal in software engineering and software development, that that needs to be independent. It's a problem because it can be very difficult to attract the right kind of people into the quality assurance role. That's probably one of the more difficult things, and why there exists these problems, is getting the right people, the battle tested people to want to become quality assurance.

It's not as sexy. The money isn't as good. It can be thought of as boring, but it's such an important role that one possible idea would be to have independent quality assurance groups within firms mandate a

0089

securities registration, for example, because that gives them both the compliance perspective, a market functioning perspective, and means that they should be paid a little more money.

So you can start to entice the right people into those independent groups.

MR. R. COOK: I just want to make sure we hear from --

MR. RIGG: No, I was just going to say I think that as most people have said here, I would see that across our clients, most clients certainly in this industry and other industries that have similar characteristics, the ability to have testers who have the right level of skill to effectively play their role is a constant challenge in this space.

Obviously in some other industries where you have a longer development cycle and you can rely upon higher degrees of documentation, more sets of tools, then you see independent testing groups or all sorts of testing groups, all things that can be effectively used.

But in this industry it's very rare, but I think the challenge though is then as we've seen in some of these scenarios, is then who in the chain of command is involved in the decision to take something into production that doesn't have an interest in that code

0090

making it into production. So often times obviously you've got the business and you've got technology, and they're working together. But it's in both of those groups' interest to move the code forward because ultimately they both have the same desire, which is to get that piece of code into production because it usually represents some feature that they anticipate is going to

increase their market share, drive revenue, et cetera.

So creating that independent check is one of the challenges.

COMMISSIONER GALLAGHER: So you're talking about the risks sort of intra-firm, whether it's an exchange, an ATS, broker-dealer. So just to shake it up, seeing as you guys are all about the secret sauce of what you do, what role could peer review play in this space?

MR. ARYA: I'd say that I think before we go to other firms for peer review, just real quick on the QA side of things, I'm all for autonomy for QA. I think QA should have an autonomous role to say yes or no in release of software.

That having been said, I was really talking about having this autonomous body imbedded within development, within the guys who understand P&L and risk. So those go hand in hand.

However, the QA guys say, you know, "We've all

0091

consulted with you, but we don't think this should go." They should have the final say.

So peer reviews, I think one of the best practices that we follow quite a bit is peer code reviews within the firm, getting guys even from other teams or within the team and actually having the right set of eyes reviewing the code, reviewing the practices, as Jamil said, really learning from your previous errors.

There's a huge amount of built in IP in the heads of all the people who have seen these failures, and when these developers or QA guys or business guys review your scenarios, that's huge.

As far as the industry peer review is concerned, I think there are problems like metrics and parallel testing and how to build better assimilators. I think we're better off coming together and having peer reviews of how to solve those problems vis-à-vis peer reviews of code and scenarios.

MR. PASTINA: I think peer reviews are excellent, by the way, within the firm. I think they're terrific because colleagues, technologists within a firm are very collegial that way, I find.

I also think it's healthy periodically to measure yourself with an independent outside group, and there are several capability maturity models that you can

0092

measure yourself against to see am I a one, am I a five; where do I rank on the scale? Am I getting any better? Am I being consistent in my process?

I think that's a healthy thing to do periodically.

MR. ISAACSON: I think there's just a level of independence. So you kind of start at the most granular, which is peer reviews, which we obviously do at BATS and they are invaluable. You absolutely need to do peer

reviews on anything that's material code. In fact, on check-ins, we have when you check in a piece of code, you need to put who reviewed the code.

In addition, independent QA. In order for independent QA to be effective, they have to know the markets, and they have to know your system. So for them to be outside of the firm, you have to pay them an inordinate amount of money in order for them to know your system well enough. They can probably learn the markets well enough, but it's going to take months for them to know the system.

So having them report independent of software development or technology but to operations, I think, makes a lot of sense. That's the way it's set up at BATS.

And then as far as independents outside the

0093

firm, you know, we do have ARP in place which, similar to NYSE, we've always taken ARP recommendations as mandates, not just as recommendations.

In addition to all of our internal and external audit functions where they're reviewing our software development life cycle, and you know, I pick different parts; so I believe there is independent audit in place today that could potentially be beefed up, but I think there should be independent audits at many different levels.

MR. R. COOK: I think Craig had a question he wants to get in there.

MR. LEWIS: Yes. Thank you, Robert.

So we've heard the perspectives of everybody and how they approach sort of risk mitigation activities within your particular entity. Yet your business models all require you to act in highly interdependent systems, and so I would like to ask a question that's related to how do you think about developing sort of system-wide best practices, and I'll put it in a particular context, and that is if you think about the demand for speed, it really drives creating different order types, and to a certain extent those order types underlying that reflect essentially nonlinear trading strategies.

So the question I have for you is if we're

0094

worried about complexity, where is the best place to address complexity? Would it be at the exchange where we would create a complex order type that may be harder to code, or should we push that off to the entity and essentially require them to create an algorithm that allows them to execute a nonlinear trading strategy in real time?

How do you weigh those two risks?

MR. ROSS: I actually think that it's not quite that easy. I mean, if you talk about high frequency trading or speed, the faster you go -- I mean, this makes

sense -- the less you do, right? If you want to do a lot, if you want to do a mark-off simulation, you're not going to do it in real time. It's just not going to happen.

So the faster you go, the simpler you get. So with that being the premise, I don't think these order types are driven by high frequency trading. As a matter of fact, high frequency trading actually prefers a smaller number of simpler order types, and that's historically how they prefer to function.

I see the order peripheral types of -- proliferation of order types in the market not driven by high frequency trading. I think they're driven by the market centers themselves, to some extent how market

0095

centers need to react to each other, the reactions from Reg. NMS, and you see actually some of those going away, some of those early order types that were done just after 2007, post Reg. NMS, actually disappearing now because they've come up with simpler ways to implement their Reg. NMS obligations.

So I think that whole premise is actually not correct. I think the basic premise is the faster you want to go, the simpler you need to be.

MR. NAZARALI: I'd like to add a point. I think the idea of high speed trading captures the imagination. You think about these co-located boxes. You think about the fact that the cables have to be equal length between the network and each of the boxes. You think about the high speed line between Chicago and New York.

It makes a lot of great headline news for the newspapers. If you look at most of the automated trading firms, most of our effort is spent on developing trading strategies to add liquidity to the marketplace. It's not really on getting a little bit faster.

Yes, that's a small part of it, but it's relatively small compared to the amount of publicity it gets.

MR. LAUER: I think the order type question is

0096

a very good one and an important one. It is the proliferation of all these order types and the complexity of these order types that is adding unnecessary complexity to the market, which is already an extremely complex system as it is, and like I said earlier, not very well understood even by the most advanced participants, especially at how these different complex systems interact.

Not only that, when you have complex order types, it leads to extremely complex testing scenarios, and you are not going to pick up all the things you could or should because you don't know what that actual matching engine logic is in general.

I think that order types should be revisited. There are so many of them now. There should be ample evidence as to the utility of every order type.

An Exchange should be able at this point to say well, we have had all these order types in place for a very long time, and here's all of the data that we can show to demonstrate the utility of these order types to the long term investor, and if that isn't demonstrated, maybe some of these order types could go away and we could have a greater drive towards simplicity, which would make testing a simpler exercise, which would improve investor confidence in different ways as things

0097

simplify rather than add complexity.

On the high speed trading side, while generally many high frequency traders are adding liquidity to the market, it was found that especially during times of market stress and the flash crash that they became market takers and exacerbated the problem.

It's a double edged coin and one that I don't think is as simple to pin down as it could be.

MR. ISAACSON: I just wanted to make a point on the order types. We are all well aware that order types go through the normal rule filing process and usually a long comment process.

A lot of the order types that have been introduced are a result of Reg. NMS, frankly. Price sliding order types, things like that, to avoid locking and crossing ISOs, ISOs that deal with the trade through rules. The order type is directly related to one or many regulations, whether it be Reg. NMS or a regulation before that. I think that is just the Exchanges responding in an effort to be compliant and offer for our members the tools they need to trade.

In addition, I think the vast majority of order types frankly has to do with routing strategies, Exchanges routing to each other. I can't speak for the ATSs or the market makers, but Exchange order types, we

0098

have a lot of routing order types based upon member demand. However, the typical high speed trader, HFT, that people think of, very rarely uses an Exchange routing order type because they have linkages to all the Exchanges themselves. They don't need our smart order router. It is usually someone who doesn't necessarily have connectivity everywhere.

A lot of the order types, as Jamil said, are not really focused -- and Jon said -- are not focused on high speed traders. They're focused on people that want a greater suite of functionality that maybe they don't have at their fingertips within their firms.

MR. LAUER: One particular order type that has recently been introduced is the PL Select order type. It seems the only explanation is for high speed traders.

That's an example of one that could be revisited as its only purpose is that it doesn't interact with ostensibly professional flow.

That again is an example of something that seems to be adding unnecessary complexity and leads people to trust the markets that much less.

MR. ARYA: I would say in isolation, most of the order types made sense, but going back to the testing point and integration testing, the whole suite of order types, if you will, actually presents a pretty huge

0099

challenge for us to actually test and so forth. There is always a demand for those order types when they come out, so we are required to really comply and also test.

I think the collection of ever growing order types requires a review, and as somebody else said, if we could review really what is the amount of volume and utility of those order types, I think that is wise to do.

COMMISSIONER GALLAGHER: Chris raises a really good point, which is very important to me, which is the extent to which regulations has driven some of the technology issues we're dealing with today versus market practices.

I think Dave raises an order type that is peculiar to high frequency trading. Chris, we have seen others. I remember years ago the New York Stock Exchange "do not ship" order type was right after NMS.

I think it is incumbent upon us, too, as we sit here with you, we are asking you to think deep thoughts and help us better oversee these issues.

You need to point out to us and not be shy about where we need to think about our regulations, our interpretations, our FAQs, whatever it is if we are driving this type of behavior, which isn't healthy for the markets in some way, we need to know about it. We need to address it.

0100

MR. PASTINA: That DNS order type remains the most popular order on the Exchange today.

COMMISSIONER GALLAGHER: That is the funniest rule filing I've ever seen.

MR. PASTINA: It was directly because of Reg. NMS and for people who wanted better control of their orders because they had decision making technology upstairs, so they did not want to be shipped to other markets. As Chris said, most of these new order types are in reaction to market structure changes and competition.

One of the interesting things, I think, about stability is the more change you introduce into the environment, the more opportunity there is for something to go wrong. One of the things about market structure and its continuing evolution is that as it continues to change and that change increases and grows more rapid, we

all will react to that, and all of the technologies have to be adjusted and they all have to be tested.

At some point, there has to be a deep breath as to how much change is actually being introduced into the system.

MR. ISAACSON: I just think when the SEC put out the concept release prior to the flash crash, we at BATS thought that was an excellent 73 pages of very

0101

thoughtful questions, frankly. It hasn't got a whole lot of air time because of the events of May 6 and Dodd-Frank and all those things.

We have had NMS in practice for five years. It was approved seven years ago. Now is probably a good time to raise that concept release again and to include obviously capital markets are almost 100 percent driven by technology, is there a way to simplify things.

One of the mantras at BATS is keep it simple. We believe we cannot operate efficiently over time unless we get rid of stuff. We cull stuff that is not used. As a whole market, I think we believe that as well. I think we should do it holistically, answering many of the questions that the concept release raised.

MR. BURNS: Following up on the point that Chris raised and Commissioner Gallagher pointed to, right now we do have the ARP program in place coming into the Exchanges, some of the ATSs, and think about it at least now acting in some respects as an independent or across Exchange review.

You talked about ways of beefing up that program, and I guess as it relates to testing or error prevention, are there concrete take away things we ought to be contemplating to focus ARP and its efforts in a risk oriented way.

0102

MR. ISAACSON: I have to commend the ARP staff. I believe as we became an Exchange in 2008, we have seen their review of us become more rigorous over time, and their focus has become much more intense. I have to commend the staff in that regard.

Regarding beefing up, it was not necessarily regarding ARP. It was independent or external audit review, which we have been through a pretty intense process of post-RIPO issues.

I can't really give specific things there other than I commend the efforts of the staff and the Commissioners, Gregg, especially getting the Midas system installed, understanding how the market interacts even as a stop gap until we have an audit trail. I think that is fantastic, the more technology and folks you have on staff that could come in and audit us.

As we were talking about QA, the utility of QA is only as good as the people that understand the markets and your system. ARP needs to have people that can come

in and understand our system at a very deep level.
I'll be honest, when we hire a software developer, I figure it's going to take him six months before he provides any utility to me, no matter how good he is, because there is just a lot of complexity. There is a lot to know in the market. We require developers to get licenses. It takes a lot for them to get up to speed. Just focusing the efforts of ARP on understanding the innards of not just Exchanges but ATSs as well would be very useful.
MR. ROSS: I think ARP provides a fair amount of value both to Exchanges and ATSs, and it's not unreasonable that you would want to see something like that for participants who are connected to an Exchange.
What I would say is if you're a firm and you're connected to a broker, third party, not connected to an Exchange, historically, you only send a few orders a day, probably don't need the extent of rigor as a firm that is directly connected to an Exchange, directly regulated by FINRA and the SEC, standard practice, the more systemic risk you pose to the system, the more highly you should be regulated. That seems incredibly appropriate.
That is the approach you have taken for ARP for Exchanges and it seems reasonable for broker-dealers to do something relatively similar. The one thing that we haven't talked about, and I know this is a topic of the next panel, is operations. There is very little literature in the finance market for best practices around operations.
Operations is really where the rubber hits the road, right. You know, everyone who has worked at a firm knows a good operations guy when they see them. It's kind of hard to define exactly what that is, right.
Operations, their whole job is to deal with conflict, conflict of interest. They are running a router that is routing out order flow for four desks and a bank, they see something that goes down, their ability to shut off their system and not face getting their feet held to the fire from internal P&Ls they have upset, is critical. That is a culture thing. That is a command and control thing that is absolutely important.
Even Exchanges, this is somewhat dated information, but when I worked in an Exchange, my regulatory department was really hesitant to shut firms off because of the fair access rules, right. That is a very straight ahead internal conflict, like Exchanges must provide fair access.
We didn't have any rules set about when it was appropriate or not appropriate to take action. It could be seen that the Exchange is taking action to shut off a participant of its own accord which is in violation with fair access.

Clearing up a lot of those internal conflicts is very important. My point is a really good operations person navigates those internal conflicts very, very well, and there is not a lot of best practice literature around that. I think that would be incredibly helpful.

COMMISSIONER GALLAGHER: On the ARP point, I hear you on ARP, and this is why I asked about peer review. I think whatever we do with ARP going forward, there has to be an assumption that the Government is not going to ferret it all out and figure it out. We can hire the smartest people in the world. We will never have enough, it will never be current enough to keep up with you guys.

I think pure reliance on ARP isn't going to get us there. It can really help, it can bring rigor and some sort of cross pollination of best practices. I don't want to put too many eggs in one basket.

MR. ISAACSON: Commissioner Gallagher, I would agree, ARP should not be the only defense or the only audit. There needs to be external firms as well as rigor within the firm.

I think peer review amongst different firms could be useful. However, in practice and implementation, I think it might be quite difficult for even the smartest guys at the firms to have enough context, like if I were to show up at NYSE tomorrow and Lou lets me see all the code, it is probably going to take me a few months before I can give any really intelligent feedback or recommendations on what he ought to be doing.

MR. ARYA: I think peer participation is really -- where it comes in handy is if we can identify problems. One of the questions was about the role of independent parties and testing and so forth.

I think there are common problems that we all try to solve, building simulators, market data replay's and so forth. If we can find ways and discuss metrics, that's where peer participation is hugely useful, not in reviewing code and so forth.

MR. LAUER: I think along the lines from the other perspective, the firms that are doing the trading and building the software, quality management standards are something to consider. It is sort of a nice middle ground between peer review and direct auditing by the SEC, for example.

CMM is one that has been thrown around and mentioned earlier. ISO 9000. Firms or industries that have adopted some form of ISO 9000, aerospace, chemicals, medical devices, health care, food safety, I'd like to think financial services is as critical or more to the well being of the country, and as such, we could look at any firm that does have direct market access having to

adhere to quality management standards.
I don't think they want to hear that because
0107
that is a substantial cost, and it would drive some players from the market, but the cost/benefit analysis should be rather clear, I would think.
MR. R. COOK: With that, I think we are at the end of the time for our first panel regrettably.
I want to thank the panelists for their insights and candor. We look forward to reviewing all of your comments you have submitted in the file.
We now have an hour and a half break for lunch before we start our second panel.
If you leave the building, please leave your visitor pass at the front desk, and make sure you allow enough time to get back through the security screening upon re-entry.
We will get back together and begin the second panel at 2:00. Thank you.
(A lunch recess was taken.)

A F T E R N O O N S E S S I O N
RESPONDING TO ERRORS AND MALFUNCTIONS AND MANAGING CRISES IN REAL TIME

MR. R. COOK: I am going to get us started now on our second panel of the day. Welcome back, everyone.
This panel is entitled "Responding to Errors and Malfunctions and Managing Crises in Real Time." As the title suggests, this panel will cover error response,
0108
focusing on how the market might employ independent filters, objective tests, and other real time processes or crisis management procedures to detect, limit and possibly terminate erroneous market activities when they occur.
I think one of the themes of the first panel was while we had a number of very helpful suggestions about best practices oriented towards the roll out of new software and how to minimize the risk of error, that errors will occur.
What we are really focusing on in this panel is what do we do about that. I think we are particularly interested in hearing from our panelists about how to approach that both from a technology perspective but also in terms of internal organizational control and people management, and how the people and the machines are interacting in the space.
We have a great panel lined up today, as we did this morning. I would like to first begin by asking each panelist to briefly introduce themselves, after which I will ask Dr. Markus to start us off with her perspectives for a few minutes, and then we will jump into some questions from the staff to kick off the conversation.
David, why don't we start with you?
MR. BLOOM: Good afternoon. Thank you to

0109
Chairman Schapiro and the other Commissioners and staff for inviting us to the panel.

My name is David Bloom. I'm the head of Group Technology for the Americas for UBS. UBS is present in all major financial markets worldwide with offices in more than 50 countries, about 36 percent of our workforce is in the Americas where I am responsible for our infrastructure.

We execute large volumes of equities on behalf of our customers with number two market share on NASDAQ for 2011 and number three market share on the New York Stock Exchange for 2011.

We are honored to have a seat on the panel. Thank you again for having us. I would also just point out that we were one of the signatory firms to the comment letter mentioned earlier this morning.

MR. C. COOK: Chairman Schapiro, Commissioners and SEC team, thanks for having me here. This is a really important step, I think, to elevate the importance of technology and how we manage it.

My name is Chad Cook. I am the Chief Technology Officer of Lime Brokerage. Collectively, we touch over ten percent of the U.S. equity volume more or less on a daily basis.

Lime Brokerage, we have been around for ten
0110
years, for over ten years. We specialize in high performance trading systems, and with that, we built in a lot of risk technology from day one.

I think our approach is somewhat unique in that respect and we arrive at our technology looking at it from good technology practices, as has been discussed in the previous panel, but also from a risk management approach coming from our backgrounds in mission critical systems and datacom and IT security.

My background, I have over 20 years building technology, primarily in datacom and IT security.

We apply a lot of these architectures and models for managing and operating technology to the financial sector, which has been pretty interesting and engaging for us. Thanks.

MS. EWING: Good afternoon. My name is Anna Ewing. I'd like to thank Chairman Schapiro, Commissioners, and the rest of the SEC staff, for having us here today. I'm delighted to be able to participate and really look forward to continuing the dialogue from this morning that I felt was very engaging.

I'd like to give you a little bit of my background. I'm the EVP and CIO at NASDAQ OMX. I have responsibility for the Global Technology organization, and that is all of the software development,
0111
infrastructure, technology operations, and information

security functions.

I'm also the head of the Market Technology business unit, and that business unit provides commercial technology to 70 Exchanges and markets around the world dispersed over 50 developed and emerging countries.

I've been at NASDAQ OMX for 12 years, and I have been part of the transformation of NASDAQ from a U.S. cash equities market to who we are today, which is a global diversified enterprise providing operations to 25 markets here in the U.S., and in Europe. We also have three clearing houses and five central security depository organizations that are part of our organization.

Before NASDAQ OMX, I have been in financial services, really my whole career, for 25 years. I've had the privilege of living through the evolution of this industry. I've had positions in both CIBC World Markets and at Merrill Lynch.

I would like to say in NASDAQ OMX's view, history clearly shows that technology has helped investors. Automation has increased transparency, lowered trading costs, and offered regulators powerful tools to eliminate bias and favoritism, and to detect and prevent fraud.

Today's technologies are inclusive, offering

0112

all market participants the opportunity to use these tools. Technology has been and remains best understood as a solution and not the problem as we continue today's dialogue.

Thank you.

MR. GAMBALE: Good afternoon. Let me thank Chairman Schapiro, Commissioners and Commission staff for inviting me to the panel.

I'm Albert Gambale. I'm the Chief of Applications Development for the Depository Trust and Clearing Corporation. DTCC is a key player in the industry infrastructure. We are the central securities depository for the United States. We have run a number of clearing corporations, and we are a designated critical member of infrastructure for the U.S.

Personally, I've been in the securities industry, all in IT, for 30 years. The last 15 years have been at DTCC. Thank you.

MR. JAHANI: Madam Chairman, members of the SEC, participants and members of today's SEC panel, please allow me to first of all thank you all for giving us the opportunity to participate in this roundtable.

My name is Saro Jahani. I am the CIO of Direct Edge, probably the newest and youngest Exchange in the equities market in the U.S. Direct Edge's market share

0113

basically has raised from one percent around 2007 to almost ten percent nowadays. We have basically shown more than anything that this type of rise would never

have occurred unless there were innovational and structural changes in the U.S. that has helped us to get where we are.

We are very confident that this panel is coming at the right time. It is very hopeful on behalf of Direct Edge to believe it is going to take us to points where we will be able to come up with some great ideas.

Personally, I've been in the area of information technology for almost 25 years, both in Sweden, where I grew up, and also here in the U.S.

I am hoping we will be coming up with some good ideas that will help us to innovate further. Thank you, once again.

MR. STEINBERG: Chairman Schapiro, Commissioners, SEC staff, thank you for inviting me.

My name is Lou Steinberg. I'm the Chief Technology Officer for TD Ameritrade. TD Ameritrade serves roughly 5.7 million retail investors who execute typically 300,000 to 400,000 trades a day.

We run a very un-conflicted model in that we are really focused on retail investors. We don't do a lot of high frequency trading and that kind of stuff.

0114

As CTO, I run technology platform engineering and development. I run technology operations and I run information security and fraud.

Previously, I had a software company in the operational risk management space focused on high availability, and over the last 25 years, I have built and run a number of software start-up's.

MR. R. COOK: Thank all of you. We really appreciate you taking the time today to share with us your insights.

This morning we introduced who else was around the table, so I'm not going to go through that again, other than to note we have Dawn Patterson with us now from the Office of Compliance, Inspections and Examinations, so welcome.

Also a special welcome to Andrei Kirilenko, who is the Chief Economist at the CFTC, and from the future side of the world, and we welcome him and invite him to ask questions as they might arise.

As I mentioned, we are going to ask Professor Markus to lead off. I will take the liberty of re-introducing her because some of you might not have heard her introduction this morning.

Dr. Markus is the John W. Poduska, Sr. Professor of Information and Process Management at

0115

Bentley University. She is also a research affiliate of MIT's Sloan Center for Information Systems Research.

Professor Markus' teaching, research and consulting interests include enterprise and inner organizational information systems, the unintended

consequences of information technology and risk management strategies, and IT in the mortgage industry.

She was named a Fellow for the Association for Information Systems in 2004, and received the AIS LEO Award for Exceptional Lifetime Achievement in Information Systems in 2008.

Thank you very much for joining us today, and please, why don't you share with us your perspectives?

DR. MARKUS: Good afternoon, all. Unlike Professor Leveson, I have only 35 years of experience doing research in the area of organizational information technology.

My background is in industrial engineering, organizational behavior, and information systems, which I think puts me in an interesting position to look at systems dynamics of the sort we are talking about here.

For the past 15 years, I have been studying the use of information technology in the mortgage industry. Much of that research was supported by the National Science Foundation and was conducted with colleagues, but

0116

the conclusions that I have made and will be talking about here are really just my own.

After the mortgage crisis, I wanted to learn what role, if any, information technology may have played in the crisis. I examined the ways that mortgage technology has evolved and how it may have contributed to the risk of errors.

I concluded that information technology did have a role in the crisis, and what I will first do is describe briefly what I learned about mortgage technology, and then I'll talk about what I think are the implications for trading technology.

I studied the evolution of mortgage technology, especially automated underwriting, from its introduction in 1994 until shortly before the housing bubble formed. That is a very short period of time, only about ten years.

In that space of time, automated underwriting technology went from being a discrete innovation with very small local scope to becoming a highly complex and interconnected web of systems that spread across many different organizations, and it touched almost everyone in the U.S. and many people elsewhere around the globe.

I'd like to talk about the nature of the changes that occurred in automated underwriting

0117

technology because I think it's important to understand how technology evolves and how that evolution, while bringing many benefits, can also contribute to risk of crisis.

First, the lending process was streamlined as it was automated. Technology became easier to use, and it was embedded in loan origination software that was

already being used by lenders.

Fewer inputs were required, fewer supporting documents were required, and outputs became simpler, too. Instead of being given long lists of conditions, people were simply told accept, caution, or refer this for more analysis.

Because the outputs were simpler, they became easier for people to act on, and people became more and more confident over time about the decisions produced by the system and stopped questioning the quality of those decisions.

Second, technology was extended to a greater variety of lending decisions. Automated underwriting technology was first applied to prime loans. Then it was applied to jumbo loans, Alt-A, and A- loans.

It was not long after that before technology was used to make decisions about subprime loans, high loan to value loans, low down payment loans, loans for

0118

manufactured housing, construction loans, commercial property loans, HELOCs, ARMs, option ARMs, cash out refinancing's, and reverse mortgages.

At the same time that the technology was being extended to many other types of decisions, the technology was also accepting a greater proportion of loan packages that were submitted to it.

Third, the technology was extended to a greater range of lending tasks. The technology first just evaluated borrower's creditworthiness. Then it performed automated property appraisals. Then it expanded into loan servicing. Then it expanded into mortgage securitizations.

At each expansion, the systems developed by one organization were interconnected with systems that were developed by other organizations.

At each expansion, new user groups began to interact with automated underwriting technology.

At the end, it was not just mortgage lenders, underwriters and credit agencies that were working with automated underwriting, it was also brokers, real estate agents, and many individual consumers who used on line lending websites.

Fourth, in the end to end mortgage process, which spreads across many different organizations and

0119

individuals, it became restructured as a result of all these other changes.

No individual or organization could see the whole picture. Some human oversight was eliminated from the process in part because people assumed that there were check points elsewhere. Fraudsters found ways around the controls.

It is important to understand that this evolution of technology didn't occur in just one place.

It was spread across many different organizations. Each organization independently tested its own software. When the systems were interconnected, the connections were tested.

The systems continued to evolve and they evolved rapidly and in many different ways, even after the connections were made and tested.

In summary, automated underwriting technology in the mortgage industry evolved from a simple discrete technology with few interconnections into a very complex, tightly coupled system that touched and was touched by many different people, processes, and organizations.

I believe that this speeded up the flow of mortgage lending in a way that fueled both the housing price bubble and speculation in mortgage backed securities and derivatives.

What does this all mean for trading technology? I understand that trading processes and technology differ from those in the mortgage industry in a number of ways.

Crucially, trading technology is even more complex and interconnected than mortgage technology. Consequently, I believe that technology is even more vulnerable to errors, disruptions, and crises in trading and markets than in the mortgage process. I think the events of the last few months are the new normal.

When systems are complex and tightly interconnected, I think there are about four basic strategies for preventing and limiting damage. The first is to reduce system complexity and interconnectedness by design, simplifying and de-coupling.

Second, testing systems thoroughly beforehand. Third, monitoring effectively in real time and reacting quickly to errors that arise, and then increasing capacity to recover quickly when fast reaction is not enough.

I think many of the panelists this morning and also I'm sure the panelists this afternoon and many of the people who posted comments on the website have come up with individual strategies that fall into one of those categories or another.

I think the most important point is that in order to assure a robust system that protects the interests not only of investors but all citizens, that these four strategies need to be pursued simultaneously. Just picking one strategy and focusing on that is not really going to address the problem.

Let me just briefly reiterate those strategies and mention a few of the many solutions that I've heard other people raise. The first strategy is to reduce system complexity, and some of the suggestions have been things like restricting trading of assets to one trading

system with managed gateways to other systems.

Another was to slow down and batch trades in high frequency trading, and I'm sure there are a number of other solutions that would fall into this category of reducing system complexity by de-coupling and inserting buffers.

The second strategy is to test systems beforehand, but I'd like to emphasize here that it's extremely important to do that, but testing alone inside individual organizations is not enough.

It's important to conduct tests on an end to end basis. I understand how difficult it could be to do this well, but I believe it's very important to do it.

Why is it difficult? Let's think about this.

0122

If we were to have done an end to end test of mortgage technology software, we would have had to bring together agents, brokers and borrowers, some of whom were fraudulent on the one hand, and on the other hand, we would have had to have investors in mortgage securities. It's not just the lenders and the securitizer's that are affected. Also, you would have had to include potential hackers and cyber terrorists if you were to do a complete end to end test of that system.

I believe that it is important to work very diligently at end to end testing, and I think that probably will require much greater industry-wide participation than currently occurs.

Tests must occur regularly, and I think live simulations that combine problem triggers that are known from past events, along with creative new scenario's like denial of service attacks, are needed in order to do the tests. The tests would also, I believe, have a side benefit. They would build skills in real time error detection and quick recovery. Simulation events like this are common in public safety, and I believe they have a place in the financial services industry as well.

The third strategy is to monitor and react in real time. Many of the suggestions have dealt with ways of monitoring transaction flow and developing better

0123

circuit breakers and kill switches. These are examples of that third strategy.

The fourth strategy is to improve capacity, to recover quickly, and I know that's a large part of the panel this afternoon, so I will say no more on that.

In conclusion, I believe complexity and interconnectedness of trading technology is such that errors, disruptions and crises may be the new normal.

To reduce the negative consequences of these events for investors and all citizens, I believe the four strategies should be pursued simultaneously, reducing complexity, conducting end to end tests, monitoring and reacting in real time, and increasing capacity to recover

quickly.
Thank you.
MR. R. COOK: Thank you, Professor Markus. I especially appreciate you reviewing the comments that we have received in the comment file, which remains open, and we encourage anyone who is interested in submitting their comments for consideration by the Commission to do that.
Bridging your experience and analysis in the mortgage industry to the securities industry, I think that is a nice segue to get into a little bit more granularity about how we pursue some of the four

0124

strategies you were just identifying for us.
Let me ask Gregg Berman to kick us off with the first question.
MR. BERMAN: Thanks, Robert. Much as we did with the first panel, I think we want to start off with the general topic about best practices, but in this case, best practices for identifying the errors and containing them as they happen.
I would like to modify that a bit from the questions you were given, just based on Dr. Markus' comments about the interconnectedness of the markets, and we all understand that, but I think she highlighted some points about the difference in the errors that can occur internally and those that arise because of the interconnected nature.
As you think about the best practices, maybe you can highlight a bit of the type of error checking and the type of concerns that you have for understanding what is going on if you have a problem internally versus if you have a problem that seems to be coming externally, either from a client or from a system you are hooking up to, and there might be differences in how you actually think about each of those things.
Why don't we start at the far end with Lou, if that is okay.

0125

MR. STEINBERG: Each incident that we have goes through a number of different stages. We start with a set of best practices around managing the incident itself. We happen to use ITEL as the framework, which we believe helps us organize the problem.
It also gives us a set of run books, particularly if it looks like an internally triggered incident, so we know how to deal with it. It gives us a set of pre-defined actions to take if we believe that the trigger may have come from an external source as well.
The notion of a trigger is important to us because there are actually three different components of an incident in our view.
One is you have a latent defect. The second is you have a triggering event, and the third is you have an

impact or potential impact which generates something we call "materiality."

We have put in place a number of things through ITEL to give us standard responses to pre-allocate resources, expertise, specialists, to help us de-bug and manage through the incident, to try to contain it if it wasn't contained in the design phase through segmentation, to then try to limit it.

If nothing else, to recover, to stop the damage and then recover quickly.

0126

The actual steps that we take on an externally triggered event or with an external latent defect are probably to separate ourselves from that source.

One option might be that we choose, if we see some unusual behavior with a market maker, to route to a different market maker, and we can do that selectively.

If it's internal, we will limit the damage, we will immediately reset the systems that are affected, work around the systems that were affected. In some cases, we use a technique we are building called "fast fail back," which allows us to effectively rewind our system's hardware and software to a previously known good state.

MR. BERMAN: Thanks. Others?

MR. JAHANI: ITEL was mentioned. It is quite interesting. I was actually hoping there would be a little bit of ITEL tweaking in today's discussion.

For us, at Direct Edge, we have deployed ITEL on a broader basis. While everybody wants to sort of have a crisis situation topic discussed, we are actually pointing out that everything is actually interconnected, the very same way basically preventing a problem and addressing them are also interconnected.

We can't just address them one by one. It has to be cohesively and realistically.

0127

ITEL has been a great framework for us because it actually has looked at the process, technology and people, the three components are there.

We have addressed this since a couple of years ago by sort of looking at a principle of having an adaptable model, a sustainable model, repeatable model, and demonstrative model. Those four components have helped Direct Edge to come up with a very robust incident management discipline.

The problem management discipline, incident management discipline, and change handling and all that stuff are also interconnected.

We have been able to create what we call a "Code Blue" model that really is there to help us to address our issues. In this situation, probably the most important thing to remember is actually the impact and nothing else, because we want to address the impact.

In order to address the impact, we can't just basically look at technology. We also need to have the compliance component, the regulatory component. By following a Code Blue discipline as a best practice model, we have been able to address this issue properly.

Once again, I want to just basically distinguish that addressing issues and managing crisis situations has to be focused on market impact, whereas

0128

when we are preventing, you are actually addressing the problems that deal with the root cause.

Dealing with the root cause at the preventive level and fighting against time is really the incident management process.

MR. GAMBALE: Gregg, I'll comment for DTCC on the front end of the question, which I think was along the lines of getting on top of and understanding when something is occurring.

We use a lot of real time thresholds and alerts, state alerts that come to us. It is generally done on a risk based analysis of an application, so our clearing and settlement applications in particular throw a lot of alerts and state condition changes out to the staff, the interested staff. I'm an interested staff, so while I'm sitting up here on the panel, I'll ignore them.

The point is that we are able to tell quickly when something is occurring, and frequently we can tell is it occurring internally, is it caused by something we're receiving from the industry. It lets us quickly adapt our incident management processes to address a problem.

MS. EWING: At NASDAQ OMX, we run our own markets, and I'll talk about that framework. We also provide technology and in some cases run the operations

0129

for other markets around the world through our Market Technology business.

Our best practices have to be very clear, very documented, very predictable, because of the nature of the type of services we provide, in addition to the markets that we operate.

Being a market operator, you are certainly concerned with your own particular issues, but you are also concerned about the other participants in the market, whether it's customer, another Exchange, et cetera.

We have a lot of monitoring mechanisms in place, not just within technology. We have dedicated operations, and we touched upon that briefly this morning, the nature and knowledge and expertise that's needed to do that operation, and quite frankly, it's getting increasingly more complex.

To deal with the complexity that's required and the monitoring that's involved, we are actually looking

at more instrumentation, more self healing mechanisms, more technology, to help complement the oversight that takes place within our operation centers.

In addition, we have our business operations functions as well as our market watch surveillance functions that are on a real time basis either monitoring

0130

the health of the markets or reacting to issues that are taking place with our customers and other participants.

Some of the very specific things we have done to protect the markets over the years, in particular, since the flash crash, are really bearing fruit in terms of helping to provide protection. If I take "clearly erroneous" as an example, we had a clearly erroneous function before, but we made the rules much more clear post-flash crash. That was one of the outcomes.

Now when we look at our stat's, our clearly erroneous filing's have gone down over 60 percent in the past two years, just to give you a metric to think about.

Monitoring, the nature of monitoring, who is doing the monitoring, how you are leveraging technology to augment the human interfaces, because there is a lot to monitor when you think about it.

Best practices. Certainly, we use ITEL and best practices around that.

The other thing, when you're in the line of fire, the more that can be pre-defined, the better. It's Crisis Management 101.

The Code Blue process. You have a standing list of contacts, names, technologies on a different call than the business, outreach, outreach to the regulator, outreach, in our case, to the press many times. You need

0131

to have a very structured escalation communication process in place.

When you are also dealing with the problem, it's not always apparent immediately what the root cause is of that problem. The ability to detect, remediate, and bring the problem to resolution is really critical and how that is managed. We have again people on call, roles and responsibilities, all types of skill sets are needed.

We talked a lot about software this morning, but our systems are more than software. It's the networks, the hardware, all the different mechanisms that connect it altogether, before you even leave your firewall.

When you leave your firewall, you have all the interconnectedness with the other markets and with your customers. That in a nutshell, the technology, the monitoring, the process, the crisis management, really all come together in reaction to problems as they occur.

MR. C. COOK: From our perspective, when we look at technology, risk, and the issues that come up,

the interconnected aspect that you mentioned is really important. As you look across even just internal systems, there are so many moving parts that are connected, and with things that are working at such high

0132

speed, it's really difficult often times to immediately point out where is that issue.

Is it something on the internal end or something on the network, something at the Exchange side. From our perspective, the way that we build technology is with the very product focused aspect, more akin to our background in commercial product development, where we look at what we are building and we say what is it that people that are using this will expect out of it.

It's much different to build an internal product where you have access to the code or you have immediate access to the networks, switches, routers, and things that you can immediately jump on and touch.

We treat all of our internal operations folks as well as our clients as our own customers.

When we build technology, we think of things in terms of not just what should the product do, but what shouldn't it do, and under the circumstances that we can enumerate, what are the situations and how do we want to allow people to respond to them in an effective way.

For us, being in a production environment, obviously we're focused on getting people back to operation quickly, understanding what the ramifications of those issues are.

I think a lot of it is driven by the technology

0133

and the emphasis on instrumenting that technology to make it enabled to be managed and enabled to be controlled and monitored in such a way that when those events that we haven't planned for come up, everybody can quickly identify and have all the data they need in order to identify where that problem is.

That is kind of an approach that we use on this stuff, along with all of the normal escalation procedures on the operational front.

It is very much driven from a technology as a tool kind of mechanism.

MR. BLOOM: I think all the panelists probably have similar -- I think someone said Crisis Management 101 -- segregation of duties, this is what should happen in a crisis, segregate your people doing investigations from those who are figuring out what to do.

When you are asking internal versus external, one of the biggest problems I've seen is when you don't know. Sure, once I know if it's internal or external, it's a pretty easy play book, and usually I can figure it out and the team can figure out what it is we need to do.

We try to make sure people in the role of triage, what is it our customers need from us in this

situation, how can we, by plugging into all the trouble shooting and everything that is going on, figure out what

it is we can do to still be there for our customers while all this is going on, or do we have to just take ourselves out.

I think that is a key point in the process that doesn't come out. Everyone has the documented incident but what are the trigger points, and when you don't know, how do you make those decisions and how do you sort of fulfill your obligations.

MR. BERMAN: To ask a follow up question, Chad, you talked a little bit about product find, if I can make up that word, of how you think about the technology because you have had external clients before.

That brings up an interesting point. During a monitoring crisis, at small firms and even at large firms, you need to get the head of development, the person who wrote the code, involved. They are the experts. That person may be on vacation. Maybe that person is not available.

To what extent -- I think we asked a question this morning about separation of the quality assurance from the development process to maintain the independence -- to what extent do you think about the monitoring process as a process that needs to be done, certainly integrated, but independently, where you do not have to be the head of development in order to know what to do,

in order to diagnose a problem.

MR. C. COOK: If I may, the productization of things is really important. When it comes to monitoring things, you have to put yourself in the mindset of who is going to be using it.

Often times as engineers, it is very easy to just think of yourself as being highly technical and you have the ability to have a good understanding.

It is important to kind of pull yourself back from that and say who are my constituents, who are the customers that are going to be using this, and what is their expectation and their level of expertise. That helps you tailor all of your functionality.

When we build our technology, we are really instrumenting it not for us but more for the people that are going to be operating it on a daily basis, whether that is our customers or operational staff, those are the important people that we look towards.

There is a feedback cycle where you work with them to say here is what we think your expectations are, here is what we are building for you, is this going to meet your needs, is it going to make it easy for you.

That, to me, is one of the critical aspects of how you design these types of systems rather than just looking at okay, we have the technical staff on hand,

0136
let's just escalate to the engineers where they can just jump right in.

It makes it more efficient and it allows us to do more innovation and more development work if we build for those folks to self manage.

MS. EWING: I would just like to add to that point. I couldn't agree with you more, Chad, in terms of productizing the operations aspect as much as possible.

We actually have what we call a "design to operate" philosophy, where the operations team is actually part of the design and functional requirements phases. Sometimes an idea, even if you think about a business idea that may sound very interesting, may end up being really complicated to surveil, monitor, or operate.

We actually even have input into the business concept that's being contemplated, and in addition, the instrumentation, the philosophy not just of how something is going to work, but how do you break it, and what happens when it's broken.

You have some basic questions to just bring it to the mindset up front, and not just for technical issues.

We as an SRO have obligations. We have rules that we file with the Commission that we have to adhere to. We also have a separate independent group to add to
0137
the QA process that is actually part of our regulatory arm.

This independent group does testing, not against the functional specification that is related to a system, they do testing against the rule book. It's a parallel segmented function that augments to again the ability -- in our case, it's not just a widget is broken, it's how we complied with the obligations and rules that we have to adhere to as a market.

That's just as critical to us from a testing and monitoring perspective as latency and other types of issues.

MR. STEINBERG: Gregg, just going back for a second to part of what I think was behind the question you asked, the notion that there might be only one developer to run to who understands the code when you're trying to de-bug, when you're trying to manage a crisis, the first thing we do is long before that, we establish a set of technology standards.

By working with a pre-established set of standards, I know what my technology components are going to be, and I can afford to invest in training, level one, level two, and level three support staff on those technologies, so there isn't a single person I have to run back to.
0138
At some point, we do go back to the developer,

but we try very hard, and in fact, we have metrics around the percentage of incidents that are supposed to be resolved at each level without escalation to the next.

The issue actually comes up in the period of time David mentioned where you just don't know what's going on, and you're still trying to triage it, so I wouldn't even know which experts to go get.

At that point, it's the operational processes that are being discussed around quick triage. Then a quick decision about am I going to try to fix this or am I simply going to try to leverage what Dr. Markus said was the capacity to recover quickly, and not try to fix it, try to get things up and running while we get to the underlying defect afterwards.

There is a very fast triage process we have to go to once we have decided what the technology component is. If we are going to fix it, if we are going to address it right now, then it's the notion of standards that allow us to have invested in the resources without having a single dependency.

MR. JAHANI: I would like to take a step back. Please allow me to elaborate. When we develop software, when we build a system, the component of risk is always there. You have to look at how much risk can you

0139

tolerate.

In the financial industry, the simplest type of systems that we build for trading, only to have high risk systems, you need to do something like at least 15,000 to 20,000 different types of test cases.

When you go to medium risk, you have to do 30,000 to 40,000 test cases. When you get to low risk, now you're talking about 300,000 to 400,000 different types of test cases.

Please imagine when you're actually talking about testing, which we all want to do, it is absolutely impossible because it takes so much time. It's going to take so much time to just regress the systems, so it becomes basically impossible.

In order to do so, the only way we can actually attack this crisis situation and so on is -- my apologies for the word -- "military precision." That's the only way we can really deal with our problems.

When there is a problem, a crisis situation, it is no time to realize who is the best guy to be here, who is the best developer, who is the best QA guy, who is the best compliance person.

First of all, everybody has to be hands on. That's the idea behind this Code Blue all hands on situation.

0140

The second thing is we need to have rigorous military sort of disciplines in place so that we can absolutely do the right thing for the market, for our

member firms, for the national market system.

In this situation, there is really not much time to talk about, really, it's a fire drill, we are fighting against time.

In order to address this situation, we absolutely have to start much earlier. We have to have mock testing, proper training. We have to make sure that there are instrumental testings as the back bone of this financial institution.

We cannot operate the Exchanges and financial institutions no longer as a development shop. We have to do it as a production shop. That is when ITEL comes in. That is when proactive measures basically help us to do the right things.

MR. R. COOK: Tom, did you have a question?

MR. BAYER: I'm not signaling any group that has mentioned ITEL in the past, but Queen Elizabeth, who owns ITEL, I don't think ever thought about a case where a market would have a melt down or a problem like the flash crash.

Knowing that IT people are introverted and they don't communicate well, we have talked about --

0141

MR. R. COOK: Tom, are you an IT person?

(Laughter.)

MR. BAYER: Yes. I don't communicate well in particular, but that's something else, from birth.

What I think is important is once you have the crisis and you now have a problem and you're going through the ITEL functions to verify the systems are up and running and you're identifying root cause, et cetera, how do you handle the social aspects, i.e., the people that are calling the ball and the people that would do a kill switch or people that would handle the ultimate reaction, how are you handling the interaction and communication?

You don't want to solve the problem while the patient -- if you were an emergency technician, you don't want the patient to die while you're administering help to the patient, you want to call 911 or you want to get additional help to come and solve the problem while you're working on the patient.

How does that happen in the Exchanges and the self regulated organizations and how do you handle the communication and social aspects, how do you call each other, how do you bring them in, how do you bring each other up to speed on what the issues are and how do you resolve them, in general.

0142

MS. EWING: I'll start. In our own markets as well, we provide, as I mentioned earlier, technology to other markets, so we have to have those mechanisms very highly defined and in place. I'm a big advocate of strong process, defined process. There are advocates of

ITEL and other standards. At the end of the day, the process can't overrule the decision making and the common sense that's needed. So, the people matter.

I've had the opportunity to work with the different markets, with very rigid structures, different degrees of rigidity, and in some cases, where everyone follows the right form, right process, they check the box, they told who they were supposed to, but you missed the bigger context, you missed the decision making, you have a longer term issue.

It's really critical that in your crisis management defined process the people that are involved -- as an Exchange, our first order of business beyond detecting what the issue is, is getting up and running or remediating the issue. Quite often we can tell what the root cause is right away.

You need to have the culture that not just comes together and understands how to work in that crisis environment, but you also need a culture and an open communication and dialogue with those introverts that

0143

maybe are the ones who oh, well, I did that change, maybe it has something to do with this, to have the ability and freedom to speak out without fear, without repercussion.

That culture is an absolutely intrinsic part of a successful ability to work and fire fight in the type of environment that we operate in.

MR. GAMBALE: I'd like to add one more piece to it. I agree with Anna completely.

I would point out to also address the point of the lone developer or sole developer, one thing that we do at DTCC, we have frequent exercises operating out of multiple locations, we will stand a location down for the day, and we will announce that the New York location will not handle any issues that come up. We can override that if we had to.

The routine crisis and instance that occur are handled by an alternate location, a location in Tampa or location at other sites in the United States.

We will rotate the responsibility, so on these stand down days, it's a little bit of a holiday, if you will, from crisis, but it lets the staff -- it gives a broader sense that the staff all around the United States can handle a situation because they're flying on their own pretty much on most days.

MR. C. COOK: If I may, on the engineering

0144

personality aspect of it, we are heavily an engineering shop. I understand what you mean on many levels.

I think from that perspective, it's really a focal point of how we build and what we are building for, so really when you look at the customers and when there's an issue that comes up, our goal is we want to get the customer back in operation as quickly as

possible.

It's less about who messed up, who made a bug, who had some issue, who pulled the wrong plug, and more about how do we work together to fix this and make the person or people or clients or whatever back in business in a meaningful way and very quickly, and with that obviously, all the right technology around it to support in that effort.

There is a certain transparency that we build into our culture. I agree with Anna's points. The culture itself of engineering is very important. This carries through to how you look at transparency within products. The previous panel, for those that saw it, they were talking a lot about you can't prevent technology issues. Really, it's about how do you recover from them.

In that, it's not about pointing fingers, it's really about making sure that people know bugs happen.

0145

We don't want them to happen. We are going to put all the rigorous testing and controls in place that we can, but it's still not going to solve every issue.

Instead, let's focus on being effective and proactive in how we monitor, how we can identify where the issue is so we can isolate it with the focal point of getting them back in business.

MR. BLOOM: When we are hiring people, the engineering discipline, those who are great at building your system and figuring out root cause after something happens, in my experience, they are usually the least well suited to do the triage while it's happening because they are too intrigued by what is actually happening compared to what do I need to do in the marketplace.

We have segregated operations staff where we try to hire what we call "applied engineers" as compared to the ones who want to be in the middle of it, and let both people do what they do best, and have a manager in charge who will make sure the right investigations are happening.

It follows up on your question about the original developer. That's for the ultra emergency, figure out the root cause, but typically in the heat of what's happening, that's not the person we want to go to.

MR. BAYER: Is the incident manager also the

0146

person who communicates with the senior management team and keeps them informed, and is that uniform across the panel?

MR. BLOOM: For us, that's the case.

MS. EWING: It varies depending on the incident. We have different levels of code. We have Code Blue, which is crisis. We have everyone in the organization from all the different areas.

If there is senior executive management on a

Code Blue level, the decision making is made at that level, because you have regulatory implications and other things to think about beyond the technical.

We have incidents every day. We all do. There are different types of incidents. It could be a router. It could be a hardware failure, et cetera.

We have different levels. That is more the incident manager tends to be the one that runs with the full issue.

MR. GAMBALE: If it's a big enough incident, we will run with two incident managers. One will be the business and regulatory liaison person working with the executives, the business, the regulators, and there is a technical incident manager who is working on technical with the engineers and developers.

MR. JAHANI: I just want to agree with

0147

everything that was said so far. I want to say we have to distinguish the problem management from incident management. Incident management really is focused on immediate recovery and basically mitigating the problems so we don't have an impact on the member firms and the market. The incident manager is literally the man in charge in that situation. If there is a crisis, the incident manager is in charge.

However, there is also an incident management group which basically is a broader group that takes the incident after recovery and basically passes it on to the problem management process, making sure that the problem gets basically resolved, the root cause gets addressed.

In other words, if we don't act this way, and we kind of categorize the incidents from say low risk incident to high risk incident, then we are practically killing the whole idea.

The incident manager has to be the person that is in charge. That doesn't take away the accountability that goes with CIOs, CTOs, or a chief operating officer. They are also part of the problem -- excuse me -- the problem management groups as such.

(Laughter.)

MR. JAHANI: The COO of the company, the CIO of the company, head of Compliance, they all have to be part

0148

of the incident management group and work on this together.

MR. BURNS: This is a very good dialogue about internal management and communication. What about communication externally, both to us as regulators and to your customers, market participants, who are wondering what's happening over there?

MR. JAHANI: In our case, like I said, as soon as there is a problem, the first person that is actually there and witnessing everything is the compliance officer.

As soon as we have an incident, a person from the Compliance group has to be there. During all our operational hours, we have full coverage from the Compliance group. They are documenting everything.

When the incident response basically is ready, that's when it formally gets communicated to the SEC, but while the problem is ongoing, when the incident is basically under management, in this situation, we immediately contact the SEC and let them know about the problem, so they are actually expecting a response as soon as possible.

MR. STEINBERG: In our case, the process is managed by one group consistently, the people that do the communications both internally and externally. Change is

0149

a function of the incident, not just the type of incident but the instant materiality.

Every single incident that we look at gets what we call a "materiality score" associated with it, which is a function of how many clients were affected, how long has it been happening or how long were they affected for, how critical are the systems that are involved, and what was the time of day, so open market, closed, it is very different from after hours. Based on that, different groups will have different responsibilities to communicate both internally and externally, senior management, regulators, et cetera.

MS. EWING: I would just add we have different mechanisms for communication. We have the same roles and responsibilities around the regulatory outreach. We have a group that does the outreach to the regulators. The client desk, the business operations function does outreach to customers.

We send out system status alerts when we even detect that we may be having an issue. That is something we will do even before we confirm there is an actual issue. The more outreach and communication you can do, the better, and the more real time you can do it, the better.

If there are crisis situations or if we have

0150

big known events, like a Russell re-balance, et cetera, we will set up industry-wide calls, and not just with ourselves, but with the other Exchanges. Those will be calls in place that are pre-planned that will take place.

MR. C. COOK: I'd like to offer a slightly different perspective. We are a pretty small organization. We don't have the large teams of staff that are dedicated to doing one function or the other.

What we have built into our culture, and this is probably pertinent to those folks out there that do have smaller staffs that are in a very operational capacity, we build it into our culture that part of that job even as an engineer is an operational aspect.

We find it to be very valuable. As I mentioned earlier, there is that feedback loop between what we would call our customer or constituent and how you evolve that technology in a way that works for them.

Having our engineers very much engaged by matter of course because we don't have that many people, having them actively engaged in any incident that comes up helps provide that feedback mechanism that we then use to improve how we do things.

We are focused a lot on automation and enabling our constituents to function on their own, self management, and things of that nature.

0151

That feedback loop is sort of built into how we operate. It's a little bit different than a large organization.

MR. R. COOK: Moving on to a slightly different topic, talk to us a little bit about kill switches. There has been a lot of discussion of this in the press, a lot of interest.

Help us unpack this concept a little bit, maybe peel back the layers of the onion a little bit and help us understand what are the different types of kill switches, how do you think about designing them, how do you think about what triggers them, how do you think about for those that require some human intervention, when is that the right way to go, who makes that decision.

What kind of process do you have around it, just sort of broadly help educate us around that whole concept.

MR. STEINBERG: I'll bite.

(Laughter.)

MR. STEINBERG: There are a couple of things to remember about kill switches. In our view, it's a kill switch if somebody does it to you, and it's a suicide switch if you do it to yourself. There is a big, big difference.

0152

People are going to be reluctant to systemically cut themselves off from the market. Any attempt to sort of pre-define automated thresholds that will in an automated way disconnect you feels like a suicide switch. The thresholds are going to be set very conservatively to a point where they may not actually operate when you want them to, at least not in the way you want them to. The counter to that is complex systems that are failing, can't be relied upon to manage themselves through that failure.

Again, we're going to worry a lot about automation kicking in at the wrong time, and perhaps de-stabilizing a system that shouldn't be de-stabilized. The ability to detect unusual behavior, whether it's based on volume or one of the many different parameters,

it is incredibly important to us, it's incredibly important throughout the whole system, when unusual behavior is detected, we have the ability to do select kills, which is to route away maybe to a different market maker.

It would probably make sense for the market makers and the Exchanges to have the ability to detect unusual behavior and reach out to us and ask why, and again, back to the materiality. We either have a fairly short amount of time to answer the question why is this

0153

unusual behavior happening or a little more time if it's less material. Answering the question and saying no, this is actually expected -- I'll give you an example.

If we see unusual behavior with one market maker, we may choose to route away and route to a different one. That's probably in the best interest of our clients and the market. If that other market maker were to suddenly see a spike of activity from us because we are intentionally routing to them, we wouldn't want them to activate an automated kill switch and shut us off because that would in fact de-stabilize an environment that we are trying to add stability into.

A quick conversation about if we can't control it at the micro segment level, having an Exchange or market maker come back to us and say based on the materiality, based on the level of impact that we see, your impact comment, we will give you five minutes to explain to us if this is normal or if this is unusual, or we will give you ten minutes, and then we're going to cut you off if you can't. If you can convince us, then that's fine.

We see sort of a layered approach to these things that probably makes the most sense with a human discussion tied to it.

MR. BURNS: Just one question. Maybe it

0154

doesn't translate but in the "clearly erroneous" context, we once had sort of the discretion versus very clear parameters and rules by which everyone has operated, so is this different in the kill switch environment?

If it is sort of a discretionary or ad hoc, what's going on over there, Fred, kind of situation, when no one else might be in that loop, does that create the uncertainty that drove away liquidity on May 6?

MR. STEINBERG: I'm not sure it does. First of all, it doesn't have to be discretionary in terms of triggering the conversation. That can be pre-defined and that can be well defined. That's the first thing I would say.

The second thing is that there is some discretion about is this expected, is this normal, can you explain it to me. If we can to your satisfaction, "you" being the Exchange, you being the market maker,

that's fine.

It is also not self said. The check and balance here is other than establishing the initial set of thresholds that would trigger that conversation, we're not doing "clearly erroneous" to ourselves, somebody else is saying hang on, now you have to convince me this is right, so it's not a loosey-goosey we're just checking ourselves and it looks okay.

MS. EWING: Just to respond to the "clearly erroneous" example and the stat I gave earlier, it is important for in this case the Exchange and our customers to have a very clear framework of how decisions are made.

It has to be as black and white as possible. When we went with the "clearly erroneous" and had very specific thresholds established for the different price points, et cetera, that provided a lot more clarity and facilitated that process, and certainly limit up and limit down will continue to advance that capability.

Kill switches need to happen at multiple layers. The working group we are a part of and we are very much a proponent of the concept of the kill switch, and at the Exchange level, it needs to complement other risk controls that are already in place, needs to complement what's in place for obligation to the market access rules at the broker-dealer level.

I think kill switches are important, but we need to ensure we don't think of them as the big red easy button. It's layered. It's complex. There is decision making criteria. There is that human element involved.

For example, one of the things we are thinking about is if you reached a certain threshold, regardless of what you end up using, we're talking about the peak net notional right now as the metric, you do an outreach, you make a phone call. You are at this level, is there something wrong, is this what you intended, do we need to increase your limits.

You definitely need that human dialogue to take place, but again, it has to be against some very defined set of guidelines and metrics.

CHAIRMAN SCHAPIRO: Is there time for that? It doesn't take very long in these markets, a couple of minutes, for an enormous amount of damage to be done. If a kill switch isn't automatic, there are pre-determined triggers and that shuts the trading down, but there's a conversation, are we checking ourselves, you have a lot of problems along the way, just from taking that time to do that.

MS. EWING: I think that is a challenge. I think that is one of the areas we are talking about and trying to define.

It depends on where we put the threshold. There is a point in time when we make the outreach. At

100 percent, if I can use an analogy, let's say at 70 percent, we make the outreach. If it hits 100 percent, the kill switch would probably have to trigger.

Those are some of the details we are talking about and designing them as we speak, and it's all about the unintended consequences of some of those decisions.

0157

It could be a very high volatility day, and that's natural order flow that we are dealing with, as an example.

Those are some of the considerations we have to design into the model itself. Again, we feel very strongly it has to be a layered approach. The kill switch at the Exchange level is really the last resort.

MR. C. COOK: On the kill switch side, I think, I agree with your points. It's almost like we're on the same team here.

I think of it in terms of the security world, a defense in depth kind of approach. Everybody has a different context of how they're looking at things, their angle, what their risk limits are, what their liabilities are.

The broker, the clearing firm, the DTCC, the Exchange, everybody is looking at one piece of the puzzle or several pieces of the puzzle, and I think collectively, you can arrive at the whole picture by putting them altogether, but that is obviously a pretty onerous task.

I guess from one perspective, the kill switches are almost going to need to exist at multiple locations and multiple places, but there has to be some form of normalized interaction where we can make that

0158

communication mechanism more efficient, and whether that's a protocol, a process, to talk with one another and say we have identified a problem on our end, do you see anything that would lead to a problem on your end.

The ramifications of some of these issues now, especially at the speeds at which we're operating, can go from very minor to very significant very, very quickly.

I think the layered approach absolutely makes sense, but then it's how do we tie those layers together and how can we work together as a community to formalize those relationships.

Typically, it's very difficult to do that with some firms because they are so proprietary in nature and there is a lot of special sauce and things that everybody has talked about.

I think outside the scope of the sauce, we don't need to interfere with what their techniques are, but we can at least arrive at a common view of risk and how to react to it, or at least give some sense of notification when we see issues.

MR. GAMBALE: Just to comment on the

coordination needed from the different trade sources for equities. When I came to DTCC 15 years ago, the morning report that showed me my trades that came in overnight and the day showed me half a dozen trade sources. Now, it has 50 equity trade sources.

I think there is a lot of coordination that needs to be done on a kill switch across all the various sources.

MR. R. COOK: I know there is a working group that is talking about that. What are some of the ideas that are at least floating around without necessarily needing to endorse any of them about how you deal with the fact that there is not just one Exchange, there are multiple Exchanges.

For example, if a threshold is hit at one Exchange, what kind of communication have people been talking about to other market participants for whom that might be relevant, and what are some of the concerns people have around any of that kind of coordination.

MR. JAHANI: I would like to comment on that. I think everybody generally speaking agrees kill switches are good and they should be implemented somehow. The question is whether or not we want to implement a red light or stop sign.

In order to be more specific on this, I think what we are really trying to say is we have to define the situation, the scenarios really, where kill switches should apply. If so, the clear protocol has to be properly identified, the thresholds have to be clearly defined and so on. I think if we can actually create that, it is a viable solution, and it should be used.

The human component, as mentioned here, it's very critical because we just don't want to do the wrong things, so if you're automating, obviously some of the scenarios definitely can be automated. The question is how many.

I think there has to be a work group that really works on the details and comes up with clear scenarios and clear protocols and procedures.

MR. BAYER: Do you envision that to be the equivalent of an open source kind of consortium where you would share monitoring activities, tools, and other best practices, and you could share them across your quality assurance staff, for example, so you would have an open source library that all the Exchanges and market participants could contribute?

Have you ever thought about that, for example?

MR. JAHANI: I believe this problem has to be taken back to its infancy. In other words, how we design systems. Please bear in mind many of our applications are not even properly monitorable because when we actually are developing code, in accordance with our

system development life cycles, we try to make them as efficient as possible, as fast as possible.

In order to really be able to get the right picture, in order to get the right flag, this is something happening, right, we need to start coding differently.

In order to do so, we have to practically look at the entire ecosystem.

I don't think any of our institutions nowadays should act as one single entity. We all are part of the same ecosystem and we have to see this as our ecosystem. Actually, we have to even take it back to our vendors, to our partners.

This is not any more a financial industry problem. People that sell equipment to us, people that sell code to us, everybody is part of the same ecosystem. We have to take it back to the design level and start doing things a little differently.

I think those scenarios all have to be identified and worked on. I definitely would agree it has to be a consortium of different roles and responsibilities that are incorporated there.

MS. EWING: I'd like to talk about one area that was actually talked about quite a bit this morning because I think it is very relevant here.

That is in the area of testing. We talked about the interconnectedness and ability to do our own unit testing, and we have automated test suites and tools that we use to do functional testing, non-functional testing, latency, et cetera. Again, going back to technology as part of the solution and tools are in place.

Where the complexity in testing really is involved is the fact that we have multiple markets that are interconnected, multiple participants, multiple protocols that connect into these markets.

We were born as an electronic market, so in our DNA, to be an open system and then how do you test towards being an open system. We have test symbols. We have had test symbols for at least ten years, maybe longer, where firms can come in, in production, and test their applications or test our capabilities, recertifying new companies as they join our markets.

We have a dedicated test facility that emulates production, in addition to being able to test in our production environment. We do weekend tests.

As an industry, and again, I've been with NASDAQ for 12 years and have been very much involved with a lot of industry weekend type testing that we used to do, that I'm going to stay one step short of mandatory, but firms would come in on weekends, and we would test. In those days, it was more new functionality or migration

0163
to a new platform, et cetera.
I think one of the things that we can and should do a lot more of as an industry is system-wide testing. System-wide testing that doesn't just test can we ping our DR site, right. Can we connect to an Exchange. We have to look at how we can test a brake, how we can test a brake as an industry, all these things we are talking about, including kill switch scenarios.
I think that is a working group, we have different tracks. One of the tracks we have is to look at best practices around not just the development cycle but the testing itself. We are a big advocate of that. We have been big investors in test symbols, in production test environments, and we would advocate and highly participate and coordinate with the industry on these types of tests.
I think that's the maturity level we are at now, and we need to do more of that. That's the only way when you have your mindset about how do you brake the thing, that's really where we need to focus next, as part of your ecosystem question.
COMMISSIONER WALTER: Can I ask a question to pick up on that and to bring in both this morning's discussion, some of which I had to skip, unfortunately, and this afternoon's.
0164
Given the kind of interconnectedness that exists today, both in terms of systems and perhaps even more importantly in terms of reputation and the reputation of the market function, how do you assure everybody plays at the same level?
We are all familiar with situations when confronted with should I spend my million dollars putting extra controls into my system or should I spend it on something that may bring profits in the door a little faster. Not everyone draws what we would call the right conclusion to those kinds of questions.
Is that the place for oversight that Professor Leveson talked about this morning? How do you make sure everybody is in the room? If you hold the test and some major player doesn't come, you have a hole in that system. How do you get to it's really going to be done?
MR. STEINBERG: In a complex interconnected system, there are very often key points where there is a lot of information flowing through. You can call it a "choke point" or a "concentration point."
Those are the places where you get the most efficient oversight. If you try to distribute it out to the edge, you are absolutely right, people tend to opt out or interpret differently.
As people come together at the market makers or
0165
at the Exchanges, you get an opportunity to start to

build the layered approach to controls and governance that is exactly what you're talking about.

You also lose something at the same time, which is you begin to extract what's happening at the edge, so with testing in particular, there is this notion of what the QA folks call "black box testing" versus "white box testing."

Black box says I know nothing about what's inside, I can load it, I can test it, I can see if it behaves to stimulus the way I expect it to, but I don't know what it's really doing.

White box testing says I deeply understand what's inside, and I can exercise very specific scenarios. To do white box testing, you kind of have to do that at the edge and you have to enforce that at the edge.

I'm not sure you can concentrate that in an industry-wide event simply because everybody's inside is different. Each participant is going to know how their system should behave on the inside.

It's both. It's a set of standards for individual participants to do things like white box testing, load and limit, test to fail, scenario planning, that kind of stuff, and then at the concentration points

0166

of the interconnected system, like the Exchanges, it's additional oversight, exactly what you are talking about.

It's the combination of the two, I believe.

COMMISSIONER WALTER: At some level, doesn't it require some sort of regulatory mandate, not at the micro level, but in a more principle based kind of way?

Otherwise, you end up with somebody opting out.

MR. STEINBERG: Yes. In a principle based kind of way, yes. The challenge would be if you tried to define at the micro level the activities at sort of the black box level, you would have an mismatch between --

COMMISSIONER WALTER: It has to be both principles based and enforceable.

MR. STEINBERG: Of course.

MS. EWING: We touched upon this this morning and we talked about the ARP and the ARP program. I actually forgot it was voluntary because we take it very seriously. That is again a program.

When you couple what you do with ARP, and we didn't touch upon this, this needs its own roundtable, cyber security, when you think of everything we are doing in the cyber security front, there is actually a lot of synergy and interplay between those types of programs.

I think we have some building blocks, some audits. We have a very strong internal audit function as

0167

well as we use external auditors to really test against best practices and standards.

Standards exist. We don't have to invent them.

We don't have to create them. We need to maybe choose them or have a short list and adopt them. We don't need to create necessarily new rules or new sets of standards. Again, I think the oversight through the ARP program is something that I think can be expanded.

MR. JAHANI: I'd like to say a few words about this. One thing that we all have to remember is when we are introducing new functionality, at least 99 percent gets tested internally in-house. All the Exchanges, all the financial institutions, they all want to do the right things. In other words, they do a lot of testing. They do as much testing as they think they have to.

The problem is basically when the issue gets out of control. In other words, it gets outside the four walls. This is when it becomes what we have experienced the last few years.

I think it was actually mentioned this morning by Chris Isaacson. Many of the Exchanges today basically have a Gemini environment. We all basically subscribe to the same market data, same production flow/order flow that comes in and so forth, and we constantly test that internally for our own use.

0168

My feeling is there has to be some sort of interconnected way. We have to kind of mandate the testing environment not should be used only on weekends, it should be used in parallel with our production, so that all the new type of orders that are coming in -- I know people see dollars and costs -- the fact of the matter is in the long run, this is going to become a major savings.

We are already operating in these environments. What if we connect them to each other, what if we actually test the test symbols or basically the real time order flow that we get, and we actually communicate to each other by facts? In other words, this is our execute on your new type of order. What is it you see here and so on.

That's when we actually start seeing the impact, and that's when we can actually report back to each other and help improve.

MR. BERMAN: I'd like to come back to the -- sorry, Tom.

MR. BAYER: Just real quick, one comment that I would add to the concept we're talking about. Anna, your concept of weekend testing, we could potentially start with a generic test period, then you could layer in some of the white box scenarios in a second test run, and

0169

then in a third test run, we could potentially layer in cyber security.

If it was all automated, we could then understand the test results. That may reduce the overall costs.

MS. EWING: Again, going back as a veteran who used to do a lot of weekend testing, again, it wasn't mandatory but firms had to show up because we had major changes going on.

What has happened is there are just so many venues and so many things to test. It is about creating not necessarily more work or more effort for the industry, it is hopefully doing it in a way that can streamline and collaborate and coordinate amongst each other so that we can do the type of testing that's needed. We have an industry-wide test coming up in a few weeks to look at our DR resiliency type testing.

We can do it as an industry. We know how to do it. We can be as efficient as possible. I think we can have different themes, and this is just brain storming now, different themes around what we focus on in any one particular test, so firms can come to the table and test.

Their internal systems, I absolutely get. We're not pretending to test beyond our firewalls, but we want to provide the infrastructure and the environment,

0170

and with ourselves, with our data vendors and other third parties, who need to also participate and be a part of that testing.

MR. BERMAN: I don't want to let you guys off that easy on the kill switch question. Just based on the initial comments to Robert's question on the kill switch, lots of complexities. There may be room for additional human layers, and there are certainly some pro's and con's. There is a very concrete idea that I think the working group has put out, and it's pretty straightforward, based on just the comment letter.

At each market center, yes, things are interlinked, but at each individual market center, there would be a programmatic way of the market center to say you can't trade, you have gone beyond the limit and that would be automated.

There might be a human call that goes out, but if there's no time for that call, it gets automated. If I'm calling you on the phone and saying you just hit your 100 percent limit, sorry about that, you're off, that's a very concrete idea that was proposed.

I just want to drill down into that and get some opinions. I think in Ameritrade's comment letter, Lou, you were specific. I think you said no, you would prefer there not to be that level of automation.

0171

I think in some of the other comment letters it was yes, we understand the complexity of that, and we understand that sometimes you might be shut off and it might make things worse, but on balance, yes, we do want that.

I would just like to drill down a bit into that. I think that is going to be an idea that is going

to gain a lot of interest and a lot of traction over the upcoming months.

MR. C. COOK: If I may, tying in the two topics of testing and the kill switch, the thing I always worry about when it comes to testing is the gap. What piece of the infrastructure, the environment or whatever, that can't be accounted for when you are doing all these tests even under the many tens of hundreds of thousands of cases.

To me, it comes down to there is one aspect of it which is just pure control. You can never control the operating systems you are using. You can never control the gear you're using. All these third party vendors and other libraries and things like that. At some basic level, you need almost that proverbial knife switch where you can pull the plug.

At that level, that obviously has major ramifications. I think as part of this kill switch idea,

0172

one of the most important aspects of understanding, going back again to sort of the risk management view of the world, looking at it from a technology perspective and saying what are the ramifications when we do pull that switch and cut off the flow of trading, what are the results, and what are those issues and the downside risks to the people we just cut off, and how do we work around those things.

I think we need to form a comprehensive view of what we want the kill switch to do, but also know what the results could be so that we can build into it other types of technology that can help offset some of the downside issues.

You had mentioned the suicide switch and that kind of thing. It's because you have those different perspectives of the person that is affected or the one that is actually pulling the switch to save themselves. I think there is a lot of room in there for development to figure out how do we solve some of those ramifications.

MS. EWING: I think we believe that we can implement a kill switch solution and we are advocating as part of the working group that we work towards that goal.

We probably got into a little bit of a requirements definition mode a few minutes ago because

0173

there are a lot of details to think about in terms of how it is designed, what kind of monitoring is needed, what kind of alerts, again, beyond the human outreach, from an automated perspective, what kind of measures do we need to put in place.

We have technology. Firms use risk management technology. We are a provider through F-10 that is used by broker-dealers to not just do their risk controls against a particular market or even particular asset

class, but across asset classes.

When you look at some of the controls that we are thinking about, whether it's credit, capacity type volumes that we want to think about, it goes beyond an one to one ratio, one to one venue. When we say it's a layered approach, it is because we believe it needs to be.

Having said that as an Exchange, we are working with defining the mechanisms of what the rules would be, how we calculate it, we have this peak notional value that we are talking about, but what kind of monitoring and alerts also need to be built as part of this.

That is another key part of what needs to be built as part of the solution, that firms can get real time messages back from the Exchanges of a pattern. This is artificial intelligence to some degree, starting to

0174

cull some of the trends that we are seeing, et cetera.

There are a whole bunch of things we are talking about that we can augment to the on/off switch, but those are the things I think need to be part of the consideration. Again, we are supporting the concept and advocating a solution. We are just driving to the details through our working group.

MR. BLOOM: Our firm is in support of that concept as well. Our view is at the start, everyone is talking about testing procedures and what not. The automated kill switches need to go through that same level of rigor.

One of the things we have to be careful about when we are setting those thresholds, be it at 75 percent where you get the phone call, versus 100 percent, where you're done, that we set those levels intelligently.

I think that is where a lot of the debate is going to come from because the last thing we want to see is a well intended kill switch disrupting proper market activity, the failure case there.

We have a lot of kill switches we run internally. I guess kill switches are now the trend, almost as credit limits. Customers come into us too often, guess what, you're done. We're not trading any more. This whole layered thing where if we see too much

0175

activity coming from one customer, they are tapped out.

If we see irregular activity on a particular venue, any of our operational staff can say we're going to disconnect ourselves from that venue and our smart order routers will take over from there.

What we are talking about here is an industry internal thing, another fail safe, guess what, maybe we are the ones doing it, who knows, we hope not, yet something sees on the Exchange irregular activity, you get the phone call first, if you don't respond fast enough and you have blown through the second limit, in

the interest of market safety, we're supportive of that automated thing happening, but again, we have to work our way into the automation using all the testing principles we talked about earlier so we don't disrupt valid market activity in the interest of trying to be more aggressive.

MR. STEINBERG: My only comment would be first, we are supportive of the concept of automated monitoring, automated alerting. We certainly have the same view of layers of kill switches. We run lots of rules starting with each order. Does the client have the cash. Is this position kind of unusual, et cetera. We start at the individual transaction level and then build up.

What we are talking about now is a much more aggregated, much more concentrated at a place in the

0176

system where we have a lot more visibility into like a market, how the whole system is behaving, so it makes sense to add additional controls there.

Our point of view is simply that getting the phone call first is a great idea, depending on how material the event is, you may have no time to react, in which case the phone call is essentially it.

If we put in pre-defined limits, that suggests that we can in advance figure out all the combinations and permutations and all the ways things might misbehave, and I think our fear of getting it wrong is going to lead us to artificially high limits, in which case we will have done a lot of work to not much effect.

What we are all talking about is sort of an escalation of you hit the 70 percent, you hit the 80 percent, you hit the 90 percent, based on how bad this is, we are going to give you less and less time to react.

We believe the fear of misfiring, the fear of de-stabilizing the market because of a combination of things we hadn't thought of could actually lead us to making and implementing a set of technology that doesn't get used the way it should, quite frankly.

MR. GAMBALE: Just a comment. As technologists up here on the panel, it's not difficult to implement a process that says I'm going to stop processing if I don't

0177

like something. We do that all the time with some of the output we send to the industry if we don't think it's in balance, we think something is wrong with it.

What the panel is saying is this is a commercial decision. It's a regulatory decision, to decide to stop processing. DTCC is very far down the food chain from where the kill switch would be implemented, but we would support it also.

CHAIRMAN SCHAPIRO: I would disagree with that a little bit in the sense that I don't think it's just a commercial decision. I think it's a broader market integrity and confidence in the markets decision about whether or not at some point if things are out of

control, trading gets stopped.

MR. JAHANI: I just wanted to add once again Direct Edge is very supportive of the kill switch model. We actually even signed a letter that was provided.

Although I want to add technology today is not the issue. Technologically, this is very easy. I think technology also can actually provide additional value to this by looking at the stop signs versus the red lights.

There is a lot of functionalities that can be embedded in a kill switch. First of all, identifying the right scenarios, creating the right protocols and all that stuff can be done, but more so, traveling models

0178

that could be used, et cetera.

It just doesn't need to be every time that there is a little problem you have to completely stop it. The phone call, but in reality, that could all be automated.

Generally speaking, we are very much behind it.

MR. BAYER: It takes active engagement if we were to agree to do this across the board. Everyone would have to be actively engaged, and we would have to continually revise the use cases for the metrics and the monitoring capabilities that are associated with it.

I think it would take a commitment from the industry to maintain those use cases and metrics capabilities.

CHAIRMAN SCHAPIRO: Someone this morning mentioned the idea of a dynamic kill switch. Is that a feasible idea? I'm not sure what it meant exactly, other than I assumed it evolved maybe as market conditions changed.

MR. C. COOK: If I may, to me, the whole idea of a kill switch sort of implies that it's dynamic. If you have it at different locations and it's looking at different thresholds and various feedback loops and things like that, it is a smarter kill switch, I think, than just the pure unplug the wire from it.

0179

One of the aspects of the control piece, the kill switch sort of symbolizes to me the fact that there is so much going on, and I think to your point, I think the DTCC is almost a central point that could help feed back into all this stuff because you have an authoritative view of what has been traded.

To me, it's that aspect of everything kind of evolving together. I kind of see it as this ebb and flow of risk limits are approaching but maybe it is offset somewhere else, and having that knowledge makes the kill switch dynamic. To me, it's sort of implied in its making.

MR. STEINBERG: I think that actually gets back into what Thomas was just asking a moment ago, which is what is the commitment to continuously evolve the

thresholds, which is really the dynamic component of it.
They just have to be either set at such an extreme level that we won't do silly things or they have to be graduated in such a way as to make sure they are used in the ways in which they were intended.
MR. JAHANI: I just want to add one more thing. Please bear in mind nobody today would like to really tolerate misbehavior, as an Exchange.
If you see that the member firm, just as an example, is choking up our sessions, you immediately pick
0180
up the phone and talk to them, and we ask them the questions.
It is not about only one situation. In single sort of situations where a member firm is causing a problem, I think we can easily handle that. Those single threaded situations are easy.
The problem is to look at scenarios where multiple things can go wrong, multiple things can go wrong. In other words, if the market behavior is changing and all of a sudden you see sort of an indication that something is wrong, it always doesn't need to be wrong on an overall basis.
These are the scenarios that need to be clearly defined. If this happens and that happens, then that's when a kill switch goes off. We need to look at those situations.
I think to your question of whether or not the industry is committed, at least I can say basically on behalf of Direct Edge, I think we are very committed to that.
MR. STEINBERG: I actually think there has been violent agreement around the concept of having well established processes and procedures as opposed to the informal processes and procedures that we already have today, and implementing them as part of a larger system
0181
that does include a kill switch.
I think just about everybody here has said yes, it's a good idea. It's then a matter of how you implement it.
MS. EWING: Again, going to the layered or multi-layered approach, when you look at the market access rules and having that in place and the requirements to adhere to that, if you look at that complementing what may happen at an Exchange level, we are looking at that holistically when we think about different layers.
I think complying with the market access rules and having the technology in place to do that is a key building block.
MR. BLOOM: I want to talk about automated kill switches also. It's easy for us to sort of imagine -- I think, Anna, you mentioned earlier potential artificial

intelligence, and fairly sophisticated ways to do it.
If I flash back to Dr. Leveson's comments this morning, I don't know that any of us would have wanted to be on the first jet if software was controlling the flaps, right.
You can imagine a scenario where we start with some very, very basic controls like we talked about, not ridiculous, but some very high limits, if you blow

0182

through it too quickly, fine, that's the automation, and everything else starts manual, and then like what happens with all the processes and how automated trading even came into being, right, you refine it from there.
I don't know there is a violent objection to automation. It's really the speed at which you do it.
MR. STEINBERG: And the combination of thresholds that you would have to put in place before you would start to trust the automation.
I doubt that any single factor, net notional, size of fill's, number of fill's, volume, et cetera, would make sense.
As I said, at least initially, I agree. Trying to define all of the different combinations of things that could go wrong in various ways would be complex.
I wouldn't want to be on that jet. I agree.
MR. KIRILENKO: From sort of the initial presentation that Professor Markus made and from what you said, it seems that the sort of fundamental issue is to diminish human oversight of systems, automated systems. That fundamental issue is not going to go away. There will be errors and malfunctions and those sorts of things happening over time.
My question is how do we assign liability for errors and malfunctions? Who is ultimately liable? Is

0183

it the programmer who put in the bug? Is it the firm? Is it the Exchange? Where does the liability rest? Should there be clear delineation of where that is?
How do we provision for that, aside from testing, monitoring, kill switches, all of that. Something will still happen. Who is liable and what is the result of that, what is the outcome.
Lou, you mentioned several times that the words "risk" and "value impacts," presumably it's in the back of people's minds when that happens. There are events that brought down companies very quickly, which became liable for a few lines of code.
What are your thoughts on this?
MR. STEINBERG: My first thought is it's not a technology problem. That is a business and legal question. I'm not sure I'm qualified to answer it.
Again, there is probably more than one thing that has to go wrong in order to have an incident. There are latent defects. There are triggering events. There

are a combination of things. The easy stuff, we have already solved for. In fact, we probably designed around it.

The concept of assigning liability is a little worrisome to me as a technologist in part because of the cultural stuff we talked about before and the importance

0184

that people not operate in a zone of fear, that technologists not operate in a zone of fear when they are trying to fix something. You don't want them hiding.

In terms of how you account for that, I'm not sure I'm the right guy.

MR. R. COOK: I think it is a very good question, but in defense of our panelists, they were sort of all vetted for their technology background and experiences and being able to describe technology to non-technologists, not necessarily to address legal issues.

CHAIRMAN SCHAPIRO: If I could go back to kill switches, what I'm concerned about is I get the idea we want a dynamic that maybe if something is going on here that changes the limits here, and we have lots of scenarios and lots of possibilities, it feels like we're adding a lot of complexity.

I wonder again if we are going to end up going down the path that we talked about this morning where too much complexity might make it more sophisticated, but it's going to make it harder to build, harder to manage, harder for everybody to know what are the rules of the road, what are the points at which their trading could be cut off. You surely wouldn't want that to be a surprise, I don't think, if you're running a firm.

I wonder if you have any thoughts about that,

0185

that we could probably make it really sophisticated and really specialized, but then have we really advanced the ball, or is a little bit blunter but simpler a better way to go.

MS. EWING: I will start. Again, I'm advocating the view of an Exchange, and I will say an Exchange that supports this.

That is a key concern. With that complexity, it leads to grayness in decision making. The ability to second guess a decision that is made.

That is that framework for decision making, that framework for what the rules are and what the thresholds are, really, really key.

We would like it simpler for that very reason. In the light of day, you need to stand behind the decision to hit that switch, whether it was automated or manual.

We couldn't agree with you more that we want to make it as least complex as possible and very clear.

COMMISSIONER WALTER: Anna, would you also be in favor, if there were sort of simple parameters set,

which of course wouldn't be right for all scenarios, if you allowed the discretion to be exercised to un-kill and to revive in a sense, to un-do the switch for that very rare case where you really think the world is going to

0186

come to an end if this switch goes off in an automated fashion?

MS. EWING: The concept about the outreach or the early warning signals is to ensure that this is a legitimate issue, especially in high moving volatile markets where you may want to change your credit limit, et cetera, because your threshold is too low, so that becomes a key element of the design of the system.

I think how you recover back from hitting the switch is something that I think would be very difficult to do, and what are the unintended outcomes and the downstream impacts, especially to other participants on the other side of that series of trades.

Those are some of the things that again we would like to simplify as much as we can, understanding it's an important decision. That is why we believe it is a last resort. That is why we believe it has to be a multi-layered approach.

I would like to also spend time thinking about how we can augment what is happening at the broker-dealer level, market access rules, 2.0, however we want to think about it.

That has to complement what we are looking to do at the Exchange level with the kill switch.

MR. STEINBERG: To your point, complexity is

0187

the enemy of availability. The only way to meaningfully automate in my view would be to in advance define all of the scenarios, which I think would drive a level of complexity.

I think baby steps make sense. I think we could have some very simple blunter thresholds. We have been talking about some of them here. I think we could leverage those in a way that sort of ratchets up how fast we need to react without trying to pre-define every permutation of every failure mode and then script that to a point where we have a high degree of complexity.

MR. C. COOK: From my perspective, it's kind of how we looked at defining functional requirements for any sort of engineering technology.

With the kill switch, there are so many different problems that could occur. We kind of have to figure out what are we trying to accomplish with that kill switch, and is it technology going awry and becoming out of control, and that's the only way to shut off trading, or is it risk limits and other thresholds or a combination.

I think through some sort of investigation where we really identify the core problems we're trying

to solve first, with that kill switch, and then some of those other elements will back off, okay, maybe there are

0188

some new regulations or principles that we can define that will surround the controls of our technology at the broker level, at the clearing level, at the other layers, so we can say these are some standards, how do you control your technology when it goes crazy or you lose sight of what's going on. How do you manage those things.

When I've been through numerous audits and other things just in terms of the normal course of business, and more recently on the Series 99 exam, which some of us may have been involved with, the gap that I see is this step from going from more human operational aspects to using technology as the core.

I always look at some of these audit things as okay, we can define the controls of who has access to the code and who rolls things out, but not once did it ask what happens when the system goes crazy and you lose access to it.

These are some of the things that are kind of low hanging fruit, I think, that if we put some of that out there as clearly defined principles or best practices or what not, how do we manage our technology now that it's such a critical part of everything that we do.

That is a little bit of how I look at it.

COMMISSIONER WALTER: Doesn't that to a certain

0189

extent go against what in some sense is the core lesson or one of the core lessons of the financial crisis, which is the thing that is going to go wrong is the thing you don't anticipate.

If you have to figure out what the scenarios are in order to define it, in some sense, I think it will fail. I think one of the advantages of a blunter tool is that it may not be the right thing in a few instances, but at least it confines the damage and doesn't require you to identify first what's going to go wrong.

MR. C. COOK: I agree. I think the risk management approach would look at it in terms of what are the ramifications we are trying to solve first, not necessarily is there trading going crazy, but what are we trying to prevent.

With those kill switches, are we really trying to prevent massive market upheaval or are we trying to prevent the broker-dealer from losing all their money because something went wrong.

I think looking at it from that perspective gives a better basis as opposed to defining every single idea. There are a million test cases that we could really look for, and that sort of goes back to one of the statements we were talking about earlier in terms of how we approach risk in the systems from a technology

0190
perspective, and what should it do, what shouldn't it do.
We can't enumerate all of those cases where things will go wrong. There are too many interdependencies between operating systems and intel CPUs, libraries, and things of that nature.
It is how do we identify when something goes wrong, and then what are we going to do about it. Kind of in that realm.
MR. STEINBERG: Just to add to that, the QA is really designed to test cases, really designed to identify and mitigate your latent defects, and operational processes are designed to be controls around triggering events.
To your point, you still will have unforeseen incidents that neither of those captured, where the only thing you can do is mitigate the impact, in this notion of latent defects trigger an impact.
What do we do when something bad still happens, and kill switches are one form of impact mitigation. There are other forms as well, since a lot of incidents are triggered by change.
I mentioned before we are implementing a model that allows us to stage changes and then if for some reason they are not behaving the way we expected, to fail back to a set of systems that have not been touched, and
0191
in fact, give us a fast fail back capability without trying to figure out what went wrong, just recover, mitigate and control the impact that way.
Kill switches are certainly one way to mitigate impact when QA didn't find the latent defects, when intel didn't prevent the trigger, but there are probably other ways that are worth investigating as well.
MR. R. COOK: We are getting near the end of our time. I just wanted to ask one quick wrap up question. Are there any other issues you think we should be thinking about in terms of reducing risk to the system.
I know some of the comment letters talked about the settlement cycle, for example, or locked in trades.
I will open it up to anyone. It is something we haven't talked about yet. We don't have time to develop it, but at least to have it on our radar screen.
MR. GAMBALE: One of the points earlier was we do think DTCC receives the most trades in real time, although not all, and we think it is important for risk mitigation and flow of the industry for us to be able to see all the trades as quickly in real time as possible.
This would help us, as Chad pointed out earlier, if we could be a participant in the kill switch feedback to the Exchanges and trade sources.
0192
MS. EWING: I would just recap really quickly

some of what was discussed this morning and earlier, and that is the real time nature of our markets, and the more we can do real time surveillance, monitoring, kill switch decisions, that is the world we live in.

As we think about the people, the processes, the technology, we have to evolve on all three fronts.

I think the collaboration that we have as an industry, this is an example, a latest example of us coming together, the working group, and advancing the idea of the kill switches, et cetera.

I have full confidence and we are fully supporting industry collaboration, and that collaboration is absolutely key to anything we advance, including with the SEC, our regulator, and having those dialogues, just like we do in information security.

There is a real tight community between the regulators, the agencies, other financial institutions, and we come together in formal mechanisms. We share information. There is no such thing as competition between us because we all have a common goal, and that is resiliency and integrity of our markets.

That is a model that is in place today that we can continue to build upon, and as we look at technology and we look at some of the mechanisms we can improve

0193

upon, whether it's the monitoring or architecture, et cetera, we are looking at the commingling of what we have advanced. There are a lot of great tools in the information security space that we are now deploying into our market systems and into the rest of our architecture.

That is another area that we are advancing. We are working with thought leaders. We are working with the Software Engineering Institute at Carnegie-Mellon, to really advance some of those concepts around how to monitor and create resilient systems through software, through controls, through more technology. I would just add that as a consideration.

MR. R. COOK: Thank you. Again, thank you all for participating today. We really appreciate you spending the afternoon with us and sharing your thoughts and experiences.

You have given us a lot to think about. I will reiterate that our comment file remains open. Anyone is free to submit their recommendations to us either building on anything that was presented at one of these panels or other ideas folks may have.

Thank you very much for your participation, and that is the end of our roundtable.

(Whereupon, at 4:01 p.m., the roundtable was concluded.)

0194

PROOFREADER'S CERTIFICATE

In The Matter of: TECHNOLOGY AND TRADING: PROMOTING

STABILITY IN TODAY'S MARKETS
File Number: OS-4-652
Date: Tuesday, October 2, 2012
Location: Washington, D.C.

This is to certify that I, Donna S. Raya, (the undersigned), do hereby swear and affirm that the attached proceedings before the U.S. Securities and Exchange Commission were held according to the record and that this is the original, complete, true and accurate transcript that has been compared to the reporting or recording accomplished at the hearing.

_____ _____
(Proofreader's Name) (Date)

0195

REPORTER'S CERTIFICATE

I, Jon Hundley, reporter, hereby certify that the foregoing transcript of 194 pages is a complete, true and accurate transcript of the testimony indicated, held on October 2, 2012, at Washington, D.C. in the matter of: TECHNOLOGY AND TRADING: PROMOTING STABILITY IN TODAY'S MARKETS.

I further certify that this proceeding was recorded by me, and that the foregoing transcript has been prepared under my direction.

Date:_____
Official Reporter:_____
Diversified Reporting Services, Inc.