U.S. COMMODITY FUTURES TRADING COMMISSION (CFTC)

TECHNOLOGY ADVISORY COMMITTEE (TAC)

Wednesday, March 22, 2023
12:00 p.m.

Commodity Futures Trading Commission (CFTC)
Three Lafayette Centre
1155 21st Street, NW
Washington, D.C. 20581

ATTENDEES:

COMMISSIONERS:

CHRISTY GOLDSMITH-ROMERO, Sponsor, Technology Advisory Committee

KRISTIN JOHNSON, Sponsor, Market Risk Advisory Committee

SUMMER MERSINGER, Sponsor, Energy & Environmental Markets Advisory Committee

CAROLINE PHAM, Sponsor, Global Markets Advisory Committee

STAFF:

ANTHONY BIAGIOLI, Designated Federal Officer

JOE CISEWSKI, Chief of Staff and Senior Counsel

PHIL RAIMONDI, Senior Counsel and Policy Advisor

TECHNOLOGY ADVISORY COMMITTEE MEMBERS:

CAROLE HOUSE (Chair), Terranet Ventures Inc., Executive in Residence

ARI REDBORD (Vice Chair), TRM Labs, Head of Legal and Government Affairs

ATTENDEES:

TECHNOLOGY ADVISORY COMMITTEE (continued):

HILARY ALLEN, Professor of Law and Associate Dean for Scholarship, Washington College of Law, American University

NIKOS ANDRIKOGIANNOPOULOS, Founder and Chief Executive Officer, Metrika

DAN AWREY, Professor of Law, Cornell Law School

CHRISTIAN CATALINI, Co-Founder and Chief Strategy Officer, Lightspark

TODD CONKLIN, Deputy Assessment Secretary of the Treasury for Office of Cybersecurity and Critical Infrastructure Protection, U.S. Department of Treasury

JONAH CRANE, PARTNER, Klaros Group

SUNIL CUTINHO, Chief Information Officer, CME Group

CANTRELL DUMAS, Director, Derivatives Policy, Better Markets, Inc.

TIMOTHY GALLAGHER, Managing Director, Cyber Risk and Investigations, Kroll

MICHAEL GREENWALD, Global Lead, Digital Assets and Financial Innovation, Amazon Web Services

ATTENDEES:

TECHNOLOGY ADVISORY COMMITTEE (continued):

DAN GUIDO, Co-Founder and Chief Executive Officer, Trail of Bits

EMIN GUN SIRER, Founder and Chief Executive Officer, Ava Labs

JILL GUNTER, Chief Strategy Officer, Espresso Systems

STANLEY GUZIK, Chief Technology and Innovation Officer, S&P Global Commodity Insights

JENNIFER ILKIW, President, ICE Futures U.S.

KAVITA JAIN, Deputy Associate Director, Innovation Policy, Board of Governors of the Federal Reserve System

BEN MILNE, Founder and Chief Executive Officer, Brale

JOHN PALMER, President, Cboe Digital, Cboe Global Markets, Inc.

MICHAEL PANFIL, Senior Director, Lead Counsel, Climate Risk and Clean Power, Environmental Defense Fund

ATTENDEES:

TECHNOLOGY ADVISORY COMMITTEE (continued)

FRANCESCA ROSSI, IBM Fellow and AI Ethics Global Leader, IBM

JOE SALUZZI, Co-Founder, Partner, and Co-Head of Equity Trading, Themis Trading, LLC

MICHAEL SHAULOV, Co-Founder and Chief Executive Officer, Fireblocks

JUSTIN SLAUGHTER, Policy Director, Paradigm

TODD SMITH, Director of Centralized Data Science and Analytics, National Futures Association

STEVE SUPPAN, Senior Policy Analyst, Institute for Agriculture and Trade Policy

COREY THEN, Vice President of Global Policy, Circle

NICOL TURNER LEE, Senior Fellow and Director, Governance Studies, Center for Technology Innovation, The Brookings Institution

ADAM ZARAZINSKI, Chief Executive Officer, Inca Digital

JEFFERY ZHANG, Assessment Professor of Law, University of Michigan Law School

AGENDA

AGENDA

(continued)

AGENDA

(continued)

AGENDA

(continued)

P R O C E E D I N G S

MR. BIAGIOLI: Good morning, everyone. I'm Tony Biagioli. As the TAC Designated Federal Officer, it is my pleasure to call this meeting to order. Thank you so much to all of our members and our non-member presenters for being here today. Before we begin this morning's discussion, I would like to turn to Commissioner Christy Goldsmith Romero, the TAC sponsor, for the welcome and opening remarks. So, Commissioner Goldsmith Romero, I turn it over to you.

COMMISSIONER GOLDSMITH ROMERO: It's so exciting to have everyone here. We've been putting this together for so long, and we just took such pleasure in trying to get the best group of thinkers, and builders, and doers that we could get in the room to help advise the Commission, and I'm thrilled that you're here today.

With the cusp of -- with our nation at the cusp of some very exciting and also challenging technological innovations, it really will take a broad representation of stakeholder perspectives to build a safe financial system, one that harnesses the best of technology while protecting customers and financial stability. And as the Commission and others are making policy decisions on next-generation technology, it's critical that we have a foundational understanding of the technology and its implications for finance and law. And that's why we have assembled well-respected technology experts for the Technology Advisory Committee.

For many of you members, this will be your first time working with the CFTC and our mission to promote market integrity, vibrancy, and resilience, and that includes instituting safeguards that make responsible innovation possible. We can greatly benefit from your expertise in determining how to ensure our markets are resilient to cyberattacks, to ensure sure that any development of digital assets protects customers, and market integrity, and financial stability, and to consider how AI, and cloud technology, and other emerging technologies can be responsibly used.

I'm exceptionally pleased to introduce TAC's chair, Carole House of Terranet Ventures, who many of you know from her work at the White House National

Security Council as the director for cybersecurity and secure digital innovation. Among Carole's many accomplishments is authoring the executive order on ensuring responsible development of digital assets. I'm also pleased to introduce vice chair, Ari Redbord of TRM Labs, who is also well known for his service at the Department of Justice and at Treasury. I also want to give my thanks to Tony Biagioli, Joe Cisewski, Phil Raimondi, LaTonya Williams, and the CFTC staff.

Today we have a panel on responsible AI, so let me start with an explanation of what responsible AI means for financial markets. In the context of financial markets, responsible AI involves using AI technologies to improve the efficiency, accuracy, and transparency of financial systems while also ensuring that these technologies are designed and deployed in a way that aligns with the interests of all stakeholders, including investors, customers, and regulators. One key aspect of responsible AI in financial markets is ensuring that AI algorithms are transparent and explainable. This means that the logic and decision-making processes behind AI-driven investment strategies and risk strategies must be easily understandable and auditable by humans. It also means that the data used to train these algorithms must be diverse, unbiased, and representative of the populations they serve.

Another important aspect of responsible AI in financial markets is ensuring that AI technologies are used in a way that minimizes the potential for harm to individuals and communities. This includes guarding against fraud and market manipulation, protecting personal and financial data privacy, and ensuring that AI algorithms do not reinforce or exacerbate existing inequalities and biases in the financial system. Overall, responsible AI in financial markets involves balancing the potential benefits of AI technologies with the need for ethical and transparent decision making, regulatory compliance, and social responsibility.

Now, I have a confession. That explanation was written word for word by ChatGPT, and it seems pretty spot on.

(Laughter.)

COMMISSIONER GOLDSMITH ROMERO: With AI being increasingly deployed in our financial system, we're pleased to hear from experts in AI from the White House Office of Science and Technology Policy, IBM, and Kroll. We also look forward to TAC's deep dive on the rapidly-growing decentralized finance -- DeFi -- ecosystem. As regulators and Congress make policy decisions related to DeFi, it is important to have a common foundation and understanding how DeFi works, how decentralized exchanges -- DEX -- or other DeFi protocols differ from centralized exchanges. For example, what are the indicators of decentralization and what they may be, how to assess the implications for finance and law. And while DeFi may hold the promise of avoiding some vulnerabilities of centralized exchanges and may hold the promise of making our financial system more accessible and inclusive, DeFi presents unique challenges, which we will hear about today.

One foundational issue is accountability. Some say that accountability rests in code, protocol, and smart contracts, or in evolving governance structures.

However, organizations may also have varying degrees and areas of centralization that could lead to accountability. I also hope there's agreement on the need to prevent illicit finance from money laundering, terrorist financing, and sanctions evasion. This is where the issues of digital identity come into play, and there are concerns also about cyber vulnerabilities. Today we are pleased to hear from several TAC members about DeFi, including our TAC chair, vice chair, the chief strategy officer from Espresso Systems, and the CEOs of Metrika, Fireblocks, and Trail of Bits.

We also look forward to the panel on cyber resilience. The new cyber strategy -- National Cybersecurity Strategy defined "resilient" as "where cyber incidents and errors have little widespread or lasting impact." And the strategy states, "A single person's momentary lapse in judgment, use of an outdated password, or errant click on a suspicious link should not have national security consequences."

Cyber resilience requires planning and preparedness so that organizations are cyber secure by design. Cyber resilience requires governance not only from the CISO's office but also the rest of the C-Suite, and cyber resilience requires reducing vulnerabilities internally, such as zero-day or end-of-day vulnerabilities, and externally with supply chain or other third-party vendors.

Today we'll hear from Kevin Stine about NIST Cybersecurity Framework. Todd Conklin of Treasury will present on cyber incident response, including lessons learned from the ransomware attack on ION markets. Deputy Assistant Secretary Conklin will also present on the benefits and challenges of cloud services technology. This is a very timely topic for our markets as critical infrastructure is considering cloud migration.

I am very honored to sponsor this tremendous group on TAC, and I very much thank you for your service.

MR. BIAGIOLI: Thank you, Commissioner Goldsmith Romero. We will now hear opening remarks from Commissioner Johnson.

COMMISSIONER JOHNSON: Good afternoon. It's a pleasure to be here for the inaugural meeting of the Technology Advisory Committee under Commissioner Goldsmith Romero's sponsorship. The work of the committee's -- of the Commission's advisory committees is critical to the -- to the development of the CFTC's regulations and policies as well as industry best practices. I thank Commissioner Goldsmith Romero and Anthony Biagioli, TAC's Designated Federal Officer, for bringing us together today. I'm also very grateful to each of you that you have volunteered your time and talent in support of the Commission's mission.

In the spring of 2000, over 20 years ago, the TAC held its inaugural meeting. The members of TAC included the chief executive officers of the largest -- of several of the largest clearinghouses and exchanges in global futures and derivatives markets. The then sponsor of the committee outlined the following agenda items, which I'm sure you'll find entertaining: oversight of electronic order routing and equation systems, common trading platforms, and common clearing. Feeling a bit antiquated today. A year later, though, following the tragic events of September 11th, the members of TAC convened at the Federal Reserve in Chicago and dedicated themselves to the tailored mission of the committee. Responding to international crises and financial markets, they steeled their focus on electronic order routing and disaster recovery, business continuity plans, and technology-centered recovery and resilience planning.

Over the last several years and two decades, TAC has continued to focus on unique and important issues at the intersection of the integration of technology and finance. Specifically, in 2005, TAC examined what constitutes prior art in the patents process, intellectual property and trading and settlements technology, restrictions on the usage of exchange settlement prices, market data privacy, and then later, high-frequency trading, algorithmic trading practices, and the role of technology and pre- and post- trading transparency as we implemented in the Dodd-Frank Act. I could go on and describe the role of TAC in advancing the conversation around legal entity identifiers, standardization of machine-readable legal contracts, data storage and retrieval, pre- and post-trade functionality, direct access market controls, and technology implementing trade execution processing and records management requirements of the Dodd-Frank Act. In other words, you all are stepping into very big shoes.

As we gather today, we consider how the world has changed. Much has been made and publicized about distributed digital ledger technology within the context of tokens, currencies, and other stores of values or mediums of exchange. Yet even if Satoshi Nakamoto's white paper, published over a decade ago, offers a precis of the archetype use case, there is much more to explore and discover in the context of the introduction of this technology in our society. Let me briefly in my remaining two minutes highlight a few.

Perhaps one of the best places to start is the remit of the CFTC and thinking carefully about the nexus that our markets have with agricultural markets. An area that I'm thoughtful about hearing from the Technology Advisory Committee on is the integration of distributed digital ledger technology in common business practices across a number of businesses in our society. For example, IBM recently developed the Food Trust Program, and in a very really -- in a very thoughtful paper, members of the Fed and other co-authors explored the distributed digital ledger technology role in addressing and reducing carbon emissions in our markets. There are any number of use cases that we could turn to and point out where DLT is helping farmers and others face challenges in data management and operations, in tracking in the context of supply chains, and answering questions regarding the verification of the source of various commodities in our society.

Another important use, which I'm very excited to hear from the committee just talking about today, is digital identity. Just two weeks ago, I spent an entire dinner conversation sitting with members of the City Corporation of London. And in our conversation, the entire focus was digital identity and the reality that in Europe, regulators are already moving far ahead in the construct and development of regulation with respect to the use of digital identities.

I'd quickly shift to a few other topics that I expect that you all will cover today that I'm excited to hear about. One of those is the risks that cybersecurity poses in our society. As a legal academic, maybe about 10 years ago, I began to research and explore the role that NIST standards play or should play in the development of business practices for the market participants we often described as intermediaries -- some would describe them as systemically important intermediaries -- and the fact that cybersecurity has an ever-evolving reality, the necessity of thinking carefully for all businesses, especially those that are part of the critical infrastructure of financial markets about how best to address cyber threats.

A few weeks ago in this room, the Market Risk Advisory Committee met and had thoughtful conversations that I believe are just the beginning, a precis, to a broader conversation that will continue today regarding how best to approach cyber threats in our markets. I look forward to hearing from the panel today and forward to -- and look forward to thinking carefully about some of the ideas that you should present.

Finally, I'd say one quick word about AI-enabled enabled cyber risks in our society, or maybe AI more broadly, and I'd share just one thought. I came to this role after having had a pretty varied career. I spent time as a practicing lawyer at a very large white-shoe law firm in New York City, and I also worked in-house for a large financial institution, but I've also had the great privilege and pleasure the last 10 years of being a legal academic. And one of the very last projects that I was working on ahead of my nomination was a book entitled, "The Ethical Implications of Introducing AI in Our Society." This book begins to explore a few of the issues that I know that you will touch upon today. I was very excited to hear Commissioner Goldsmith Romero describe the necessity of thinking about transparency and explainability in AI. I join her on this soapbox. I've been there for a number of years. I'm quite excited to hear how we can think carefully collectively about the best way to mitigate replication or redundancy of discrimination through the use of certain data sets or data practices.

I'd close just by noting that I'm very grateful that I have the opportunity to serve alongside Commissioner Goldsmith Romero. She has proven to be one of the most exceptional individuals that I have had the privilege and the pleasure of working with. I'm grateful that we were nominated the same day and am excited continuously about the opportunity to work with her with, with Commissioner Mersinger who is here in the room, Commissioner Pham, and our chair, Ross Behnam, who's not with us today. Thanks so much.

MR. BIAGIOLI: Thank you, Commissioner Johnson. Commissioner Mersinger?

COMMISSIONER MERSINGER: Thank you, and good afternoon, everyone, and thank you for everyone who's here in person and those who are also joining us virtually. I'm really looking forward to today's meeting. I want to commend Commissioner Goldsmith Romero for convening the TAC Advisory Committee and for putting together such an impressive group of presenters for today. I really expect this to be a fascinating discussion.

I also want to acknowledge Tony Biagioli. I think we should all, like, keep track of how many different ways we pronounce his name today. But so he's the Designated Federal Officer for TAC. It takes a lot of work to plan and organize these meetings, and Tony was able to do this, accomplished it all while he has a day job in our Division of Enforcement. So thank you, Tony, for all your hard work. Additionally, just want to thank the CFTC staff that work behind the scenes to make sure these meetings happen, whether it's telecom, logistics, IT, security, many teams that we have involved planning and executing these meetings. We wouldn't be able to hold these meetings let alone do our job without their expertise and hard work.

Every topic on today's agenda is timely, relevant, and critically important to the American economy. As regulators, we rely on your expertise to help us do our job in a way that allows responsible innovation to flourish in the derivatives markets we regulate. Our governing statute, the Commodity Exchange Act, in it, Congress has specifically tasked the CFTC with promoting responsible innovation among derivatives markets and market participants. But while Congress directed us to help assure that our regulated markets reap the efficiencies and benefits of emerging technologies, it also requires us to do so in a manner that ensures both market integrity and customer protection. That can be a difficult balance to achieve, and we cannot make those judgments without a better understanding of those technologies. Sound policymaking comes from opportunities like today's TAC meeting where we can engage with the public, gather information, and learn from those who are most knowledgeable in the field to inform our regulatory decision making.

I appreciate all the time and effort from all of our presenters today as well as those who serve on the TAC under Commissioner Goldsmith Romero's sponsorship. Your service on this advisory committee is truly a public service. I firmly believe that government action without public input is misguided at best and, at worst, it could actually create more harm than good. That is why all five of the CFTC's advisory committees are essential to the work we do at this Agency. So again, thank you all for being here, and I really am looking forward to the presentations and discussion.

MR. BIAGIOLI: Thank you, Commissioner Mersinger. We will now hear pre-recorded opening remarks from Commissioner Pham.

COMMISSIONER PHAM: Good afternoon. Thank you to Commissioner Goldsmith Romero, Tony Biagioli, the Designated Federal Officer, and all of the members for today's meeting of the Technology Advisory Committee. Today's discussions on cybersecurity, decentralized finance, and artificial intelligence are incredibly timely and important to the mission of the CFTC. I thank Commissioner Goldsmith Romero for her leadership in tackling these issues with renowned technical experts, and I thank all the guest speakers and members who are willing to share their time and experience with us today.

In light of the ongoing shocks and disruptions to markets, I would like to focus these brief remarks on operational resilience. I think it is important to note that for many years now, both policymakers as well as the private sector have identified and recognized the vital need for operational resilience in our financial system, and I support this for both financial institutions as well as financial market infrastructures.

The Financial Stability Board, the Basel Committee on Banking Supervision, and the International Organization of Securities Commissions, and regulatory authorities around the world have done significant work to strengthen operational resilience and identify vulnerabilities. The CFTC is actively engaged in these international efforts. As noted by U.S. prudential regulators in 2020, operational resilience encompasses governance, operational risk management, business continuity management, third-party risk management, scenario analysis, secure and resilient information system management, surveillance and reporting, and cyber risk management, and I'm pleased that we are focusing on these risks today.

Among our various registered entities and registrants, the CFTC has direct oversight over both U.S. and non-U.S. global, systemically-important banks registered as swap dealers, as well as three systemically-important financial market utilities that are registered with the CFTC as derivatives clearing organizations. You can see that the CFTC has a critical role in ensuring financial stability and mitigating systemic risk.

Establishing and maintaining a robust regulatory framework to manage the risks that are part of ensuring operation resilience is core to our mission, and both the CFTC and our partner, the National Futures Association, have rules that are already on the books. Many of the recent disruptions, including ION Trading, are addressed in these regulatory requirements. Accordingly, I believe we must examine and address compliance failures under our existing rules as well as considering whether additional regulation is necessary. I look forward to hearing today's panel discussions and continuing our public engagement on these topics. Thank you.

MR. BIAGIOLI: And thanks to Commissioner Pham. Thanks to everyone for opening remarks. Before beginning our first segment, there are a few logistical items that I've been asked to mention to the committee members.

Please make sure your microphone is on when you speak. This meeting is being simultaneously webcast, and it is important that your microphone is on so that the webcast audience can hear you. If you'd like to be recognized during the discussion, please change the position of your place card so that it sits vertically on the table, or raise your hand and either Carole, or Ari, or I will recognize you and give you the floor. If you're participating virtually and would like to be recognized during the discussion or for a question, please message me within the Zoom chat, and I'll alert Carole and Ari that you'd like to speak.

Please identify yourself before you begin speaking and signal when you are done speaking. Please speak directly into your phone -- into the microphone for optimal audio quality on the webcast. Please unmute your Zoom video before you speak and mute both after you speak. Please only turn on your camera when you are engaged in discussion, and if you're disconnected from Zoom, please close your browser then enter Zoom again using the previously provided link.

Before we begin, we'd like to quickly do a roll call of the members participating virtually so we have your attendance on the record. So after I say the name of our several virtual members, please say that you're present and then mute your line.

First, Christian Catalini?

MR. CATALINI: Present.

MR. BIAGIOLI: Jill Gunter?

MS. GUNTER. Present.

MR. BIAGIOLI: Jennifer Ilkiw?

MS. ILKIW: Present.

MR. BIAGIOLI: Michael Panfil?

MR. PANFIL: Present.

MR. BIAGIOLI: And I believe that's all we have for now. Before we dive into our first topic, it is my pleasure to introduce the newly-appointed chair of the TAC, Ms. Carole House, and the newly-appointed vice chair of the TAC as well as our first presenter, Mr. Ari Redbord. Carole, I'll turn it over to you.

MS. HOUSE: Thank you, Tony. Good afternoon, everyone. I'm honored to chair this inaugural meeting of the Technology Advisory Committee to the CFTC. I would first like to thank Chair Behnam, Commissioners Johnson, Mersinger, Pham, of course Tony and CFTC staff, and especially Commissioner Goldsmith Romero for her leadership and her vision in sponsoring the TAC's reconstitution, and bringing us together all here today to discuss and advise the Commission on critical issues related to technology's impacts to financial services and commodities markets.

To help frame our discussions today, I'll underscore that technology sits at the heart of the U.S. economy and financial services. It shapes the way that institutions are providing those products and services to consumers and engaging with other players across the financial system. Technology plays a critical role in all elements of the risk equation. It shapes the conduct of those services being provided, the nature of threats and vectors attacked by illicit actors, the vulnerabilities that can be exploited through malice, negligence, or otherwise risky behaviors or conditions. And finally, technology also plays a critical role in mitigations, providing innovative capabilities, which, when implemented responsibly, can help industry, regulators, supervisors, law enforcement, and national security authorities to detect, prevent, and disrupt different kinds of risks in financial services.

Technology is a tool for licit actors and illicit actors, and it has implications for a spectrum of policy issues that all matter to the Commission and the broader regulatory and U.S. Government community, including market and operational risk and resilience, economic competitiveness, illicit finance and fraud, environmental impacts, financial inclusion, and equitable access, just to name a few.

My fellow members of the Technology Advisory Committee, Vice Chair Redbord and I are here to serve the Commission and to contribute not just an understanding of the current state of play and the challenges presented by technology but also to discuss, debate, and distill potential possible solutions and ways forward that are aimed at helping the Commission better understand and address these challenges, as Commissioner Mersinger mentioned in her opening comments, to help them drive responsible innovation.

Today we will focus our discussions around decentralized finance, cybersecurity, and responsible use of artificial intelligence, all of which are key issues affecting the current and future environment for finance, regulation, and supervision, including for commodity markets specifically. The TAC leverages an incredible scope of expertise here to inform our discussions. I am honored to be surrounded by thought leaders and experts across a variety of sectors, representing institutions and backgrounds in capital markets and trade finance, banking law and regulation, prosecution regulation and compliance for countering illicit finance, cybersecurity and data science, environmental security, ethical application of emerging technologies, venture, cloud and infrastructure services, reg tech, academia -- I could go on. This is an incredible team today, and I'm sure the whole committee will join me in thanking the Commission for this opportunity to serve.

So now I will turn to Section One of our agenda. It is my pleasure to introduce our vice chair and first presenter regarding DeFi issues, Ari -- Mr. Ari Redbord, head of legal and government affairs at TRM Labs, and will present a brief survey of the DeFi landscape.

MR. REDBORD: Thank you much. A really true honor to be here today. Thank you, Commissioners Johnson, Mersinger, and Pham, and a very special thank you to Commissioner Goldsmith Romero for your sponsorship of this committee and to my fellow committee members for your service. It is a true honor to serve beside Carole House, our chair, as she's a former Treasury colleague and friend.

I've spent my career working to protect the financial system from illicit actors, first, for over a decade as a prosecutor at the U.S. Attorney's Office for the District of Columbia, and then at the U.S. Treasury Department, and now at TRM Labs. But as we kick off the work of this committee, the focus should not only be on the risks but on the promise of the extraordinary technology that has the potential to not only change financial services, but the very ways in which we interact with each other.

Any discussion of regulation of decentralized finance should begin with the promise of decentralized finance. In the wake of the collapse of FTX, we woke up every morning to headlines like, "FTX Predictable Failings Show the Need for Crypto Regulation." That was the Financial Times. "Will the Collapse of FTX Lead to Better Crypto Regulation," from the New Yorker. However, in reality, FTX had very little to do with cryptocurrency.

As a young lawyer in the age of Enron, WorldCom, Lehman, FTX looked very similar: a case of fraud, a lack of corporate governance, and the commingling of funds. The fraud at FTX did not occur on blockchains. It occurred in the opaque quarters of centralized financial institutions. Even prior to the collapse of FTX, when we have thought about crypto policy, it has been in the context of centralized exchanges, like

FTX, with regulators seeking information from siloed intermediaries, the same way that information flows to from banks to their regulators today.

However, the true promise of blockchain technology is DeFi. DeFi is financial services offered without a traditional financial intermediary and delivered via a software program or smart contract, which uses distributed ledger technology and enables peer-to-peer transactions. DeFi enables an ecosystem of peer-to-peer financial services untethered from many of the issues that plague our current system and offers the promise of financial inclusion: peer-to-peer, cross-border value transfer at the speed of the internet. That is the promise.

DeFi allows users to access most banking services, such as earned interest, buy insurance, trade derivatives, trade assets, borrow, lend, and more, but without requiring paperwork or third-party involvement. I start with the promise of the technology because it is critical to understand what the technology enables as we discuss what policy could or should look like.

The promise of decentralized finance stems from the native properties of public blockchains: data that is transparent, traceable, public, permanent, private, and programmable, and can allow anyone, from regulators to financial integrity professionals, average citizens to law enforcement, to more readily identify risks to the financial system. I'm going to go through a few of these qualities now.

First, the data is transparent. The nature of public blockchains as open and distributed ledgers means that each transaction is verified and logged in a shared immutable record along with the timestamp of the transaction and the blockchain addresses involved. This data from the public blockchain is transparent, enabling the financial industry and government agencies to monitor trends in financial crime, market abuse, and financial stability in real time, and conduct more effective risk assessments.

But it is more than just regulation. When we talk about things like proof of reserve, which is very top of mind right now, the proof is on the blockchain. The technology to provide auditability and transparency has been inherent since inception. Data is traceable. Because blockchains provide an immutable audit trail of every transaction, understanding the ultimate source and destination of funds, particularly across jurisdictions, is substantially easier, faster, and more reliable compared to tracing funds through traditional financing mechanisms. An example is the attack on Colonial Pipeline, where a -- where a -- where a ransom payment was made in bitcoin and was then ultimately able to be tracked and traced to an address that the U.S. law enforcement authorities were able to seize back.

The data is public. Unlike transaction and customer data held by companies or financial institutions, public blockchains are distributed and not managed by a central authority. Thus, anyone, including law enforcement officials and regulators, can access, identify, and trace blockchain transactions as the information is free and publicly accessible, independent of a third party.

The data is permanent. Storing transaction records for long periods of time is costly, cumbersome, and may be prohibited under local law. In contrast, transactions are permanently recorded on the blockchain, which allows institutions, auditors, and government investigators greater ability to follow the money, even if the transaction is several years old. An extraordinary example of this is the 2016 hack of the Bitfenix exchange where the launderers ultimately spent years, and through myriad obfuscation techniques, to move funds, while law enforcement authorities were able to go back because those records were logged on an immutable public ledger, and trace and track the flow of funds, ultimately recovering the largest seizure in U.S. history five or six years later.

And finally and arguably, most importantly, the data is private. As more and more consumers, businesses, and governments transact on blockchains, it becomes even more important to enable financial privacy on blockchains in order to protect consumer privacy, prevent corporate and national -- nation-state espionage, reduce the risk of data breaches, and protect national security.

It bears emphasizing that privacy and blockchains are not incompatible. In many ways, blockchain-based technologies, by minimizing the need to store personal data in one centralized repository, by empowering individuals to assert control over who accesses their data, and by allowing individuals to determine for what purposes their data will be used, are more privacy protected than the status quo. There are extraordinary technologies being built today. I know we're going to hear from Jill Gunter and Chair Carole House a little bit later about the real promise of some of these technologies and really looking forward to that important work.

The data is programmable. Blockchain provides a new opportunity to increase access to the financial system by reducing the cost of providing financial services and programming key outcomes through smart contracts. The promise of DeFi is the technology itself. To date, the conversations around crypto policy and regulation have been about how to jam crypto into the current regulatory paradigms, how to regulate the next FTX essentially, but the native qualities of public blockchains allow for a different regulatory paradigm that balances the right to privacy with the need for security. I believe that this committee and this real extraordinary group of subject matter experts is the perfect place to begin those conversations and look forward to the conversation.

MS. HOUSE: Thank you for that presentation, Ari. For our second presentation regarding DeFi issues, Ari will jointly present with Mr. Nikos Andrikogiannopoulos, founder and CEO of Metrika, on the topic of decentralization indicators and issues. So Ari and Nikos, take it away.

MR. ANDRIKOGIANNOPOULOS: Thank you so much. I cannot think of a better time, place, and audience to be talking about the decentralization today. Starting with kind of a little bit of the history of decentralization, decentralization is not new. If we look throughout history, there has been -- even going back to ancient Greece, democracy is a great example of decentralization, entities on their own deciding for -- what's best for the future of their communities, what's better -- what's best for the future of their societies.

The decentralization in computer science became known over the past 40 years. There is the famous problem, which you can see in the picture, of the Byzantine generals where they're trying to attack a city, and there are multiple armies and multiple generals surrounding the city. And they're sending messengers with -- they're sending notes with their messengers, but they cannot trust the messengers to coordinate an attack. And that problem has fascinated computer science over 40 years. How can we coordinate between parties that we cannot trust?

This has been solved. With the advent of technology, with modern cryptography, with consensus mechanisms, this a done deal. This is a solved problem on how we do it. But when we look at, you know, the society and economy, particularly the financial vertical, this has been the least centralized kind of aspect of our economies. And I think it really brings to mind how can we have finance take a deeper look into decentralization that benefits the nuances, and how can we understand in depth, that we can -- we can adopt it, and that's what we're going to talk in the next slides.

Decentralization has many different dimensions. The classic definition of "decentralization" refers to transfer of control and decision making from a centralized entity to a distributed network. That's kind of the textbook definition. A lot of people, when they talk about decentralization, they talk a lot about the pendulum, something that keeps moving over time even if one part goes away, something that withstands the test of time.

And when we look our -- when we look through the glass into the different dimensions of decentralization, we can look at technology and the source code. There are multiple developers where each one of them can individually make their own decision where do they want to contribute, how do they write their code. They can each decide for themselves. When we look at the network, there are so many different network elements around by different operators. They can decide how to run operations best on their own. They can decide their hardware, their configurations, how they run their business.

In terms of custody, we have a variety of options where people can hold their assets. Similar to how we choose our own email client and we can make our decision, we can choose a self-hosted wallet and monitors wallet. We can make decisions individually on where we host our assets. When we look at dApps and DeFi, which are all empowered by smart contracts, those live on the chain. They're being validated by different individuals who can choose to either approve or not approve the execution of that logic, of what the application does. And last but not least, and this is where kind of the human element goes into that, those ecosystems, the community, can evolve through decision making through the economics, and they can decide how they iterate and how they evolve the governance system that they have.

All of what I just described are many different aspects of decentralization, and each one of those components is decentralized a certain degree. On top of that, on top of this micro picture, there is a macro picture of decentralization. It's not just one blockchain network. It's multiple networks connected with each other with bridges, with oracles bringing information from the outside the world in, and with on- and off-ramps. So every element that I just described has a certain degree of decentralization in it, so when we look at the broader picture, we have to take into account the level of decentralization of all of those components that make the ecosystem.

Talking about the benefits, and I believe Ari talked a little bit about that, it increases the transparency and accountability. Everybody knows the ledger is the source of truth, the undeniable source of truth, and one can wonder could things like that have prevented the SVB problem with decentralization. Could it be that if we had access to real-time proof of reserves and proof of liabilities, and everybody -- there was no information asymmetry, and everybody knew what is happening at any given point in time, could that have prevented a shock in the market?

It enhances security. Bitcoin has never been hacked. Ethereum has never been hacked. Many networks have never been hacked out there. And one can ask the question, the data breaches that we have of centralized servers, like the New Zealand Bank, the Robinhood data breach, could those have been prevented if there was decentralization built into the architecture?

And the last thing is enabling greater autonomy and control. Users have more control. They're not just customers over a vendor, but they can take on more responsibility. And if you think about businesses historically, they all try to bring their customers closer to them. They try to build some loyalty, either through airline miles, through various mechanisms. This brings that into an equal footing. The consumer now becomes partly the producer as well, and it's a kind of a more integrated relationship between vendors and customers in this ecosystem.

On the challenges side -- on the next slide on the challenges slide, obviously there is a bootstrapping problem. This technology is pretty new, so once these networks get created, once they -- until they get sufficient size and a sufficient degree of decentralization, the benefits cannot be realized, and the risks are quite significant. So there are some threshold levels that need to be achieved so that decentralization is effective.

There is also an aspect of technological maturity. These networks need to scale to be resilient, to be reliable. We cannot have networks going down for 20 hours and people not being able to move their assets. They need to be interoperable. You need to be able to kind of cross from the one network to the other so that you have options.

And when it comes to governance, and this, again, speaks to the human element involved in this, there is the tragedy of the commons. Without active participation by the community members, proposals can get rejected. Proposals can be accepted with no one really paying attention. So the way they evolve and the direction these ecosystems evolve pretty much depends on how we can have active participation, much like in a democracy. And the last thing is the interconnection of the decentralized networks to traditional finance. Traditional finance needs to manage the risk. That's what traditional finance does, so they need the right tools so that they can understand and embrace decentralization.

So all in all, the benefits of decentralization that's very well-articulated over the centuries far outweigh the challenges that I described. And some of the challenges will self-resolve as a function of time, and technological progress, and maturity. However, I think when it comes to governance, when it comes to risk management, and where the human element is involved in those newer technologies, we also need new tools. We need new practices so that DeFi can be enabled for broader adoption. And I think we've reached the point in time where we can no longer ignore decentralization. Not only we have to embrace it, but I think it's our duty to lead it in the right direction. Thank you.

MR. REDBORD: Thank you so much, Nikos. Building off of your remarks, which were -- which were terrific, and building off of sort of my opening remarks, look, DeFi enables this extraordinary ecosystem of financial services, and there are truly extraordinary companies building in this space, and DeFi matters now. The total value locked in DeFi has exploded in the past two years from about $10 billion in October 2020 to $47 billion in February 2023. DeFi was stress tested during FTX and some of the recent events and did not fail. DeFi is absolutely here to stay. But that said, there are vulnerabilities as there are in any ecosystem, and I'm going to talk to you about some of those vulnerabilities today.

The first is the technology risks, the hacks, the code exploits that have become a way too regular occurrence in the DeFi ecosystem. 2022 was a record year for hacks -- $3.7 billion in stolen funds overall in the crypto ecosystem, 80 percent against DeFi targets -- and these were the largest hacks. You see the largest hack, the Ronin Bridge hack, for over $600 million on a -- on a -- on a bridge connecting Ethereum to the Ronin blockchain. Hacks have become an everyday occurrence, and they've become more and more perpetrated by nation-state actors, like North Korea.

Frauds and scams are something that we're seeing in the DeFi ecosystem, and they seem to be getting larger and larger. We identified about 11 what we're calling mega investment fraud schemes, $100 million or more in 2022. As we build the system, it's -- as we build this ecosystem, it's so important to keep illicit actors from taking advantage of this new technology.

Sanctions, obviously something Commissioner Goldsmith Romero mentioned in her remarks and something we all need to focus on and ensure that we're hardening defenses against. How should we do sanctions compliance in a decentralized space? How do we ensure that bad actors are kept out of this new ecosystem and something that I'm hopeful that this committee can spend some time on over the course of the next several months.

Market manipulation, something we've been seeing more and more as the ecosystem grows, with an example of Mango Markets, you know, what is legal, what is not legal when it comes to market manipulation in the DeFi space, important issues that we should be thinking about when we're thinking about vulnerabilities. And finally, money laundering. We're seeing illicit actors move funds across blockchains in and out of decentralized exchanges. But one thing that's so extraordinary, and going back to my sort of initial remarks, is that transactions are visible, transparent, immutable, meaning that anyone can watch these financial flows, can trace and track the flow of funds, can share information amongst a community and, ultimately, attempt to stop this type of illicit activity.

We're seeing the development more and more of the development of privacy-enhancing technology, which looking forward to hearing more from Chair House and Jill Gunter later on today. This could be really important when it comes to identity and mitigating some of these risks. And finally, something that's really important to note and something that we're seeing as we look on chain, is that while you have these vulnerabilities within the DeFi ecosystem, all roads still lead to centralized exchanges. As we say, all roads lead to VASPs. In other words, illicit actors are still needing on-ramps and off-ramps into the DeFi world, and that's still where, obviously, you're seeing conversion. You're seeing the type of money laundering that one is worried about. So while there are obviously vulnerabilities in the DeFi ecosystem, the real vulnerability still exists in the sort of more centralized space.

And finally, there's real promise in sort of the regulatory space today. You know, as we, obviously through this -- the work of the CFTC, think about these issues, you know, the paradigm today when it comes to regulation is these siloed institutions, these intermediaries reporting directly to their regulators and, frankly, never see each other's transactions and don't really understand sort of what's happening within these walled gardens. And what really this technology enables is the ability to really think about regulation in an entirely different way where we all have, whether it's an individual or a government entity, we all have visibility on flows that -- in ways that we, frankly, never had before.

So really looking forward to continuing this conversation, and I am going to hand it over to -- I am actually not going hand over to anyone. I am going to open the floor to questions, comments from committee members, and I see signs standing on their end, so that is fantastic, yep. Oh, terrific. Stanley, why don't I -- why don't I start with you? You're the first one up.

MR. GUZIK: Great. Thank you. Well, I just want to say to all the commissioners, thank you. I really appreciate this opportunity. Just a couple of comments on, Nikos, what you were talking about.

In the DeFi world, we talk about all the benefits, and we also talk about the risks, and some of the risks that, you know, I would encourage us to consider is risks at the protocols. When you're talking about DeFi, there's, you know, the tokenization -- the tokenization of proof of work, which we know that there's a limited number of these, you know, digital assets with proof of work, but the industry -- many of these protocols are moving to proof of stake. And with proof of stake in these distributed DeFi environments, the nodes that are processing the environments, you have to stake X amount of -- X amount of tokens.

So I'll just use the example of Ethereum where Ethereum moved over in September from proof of work to prove of stake. To get a node running on the Ethereum network, you need to stake 32 weeks. But now what ends up happening is the protocol -- like we mentioned, these protocol open-source bodies, who -- it now starts becoming a centralized body for the minting of new tokens. So we moved from a -- an algorithm controlling how many tokens could be minted to centralized bodies with proof of work. So I think that -- you know, that's one of the things I would encourage this panel to discuss.

And then the other part of that is it is open -- these protocols are open. You have nodes on the networks, and, you know, what is the risk of a decentralized network now becoming centralized because, you know, the emergence of big players coming into the market who are running thousands of nodes, and I think roughly is about 400,000 of Ethereum nodes now validating, basically validators. So you could actually run the risk of these large companies coming in, setting these centralized -- you know, it's a decentralized network, but it actually becomes a centralized network with larger companies running validator nodes.

MR. REDBORD: Thank you so much. Hilary?

MS. ALLEN: Yes. Again, thank you so much for having me here. So as a basis for our discussion, I thought I'd offer a bit of an alternative perspective on DeFi. I've done a lot of research on the space, and my findings are a little bit inconsistent with some of the descriptions we've heard, so I just thought I'd offer those as an alternative or a complement to our discussion.

So I think it's important to recognize the difference between technological decentralization and economic decentralization. So most of the -- what we're seeing in the DeFi space is technological decentralization, and this relates actually to the comment just made, which is, that's all well and good, but if you have economic centralization behind it, you lose the benefits of the technological decentralization. And that, I think, is very much what we see in the DeFi space at the moment.

We see intense economic concentration, holders of governance tokens. And those are very, very concentrated, and then there's a lot of other centralized intermediaries. We heard about the oracles, the data feeds. Those are often centralized data sources, et cetera. So this space is not economically decentralized in any sense, to my mind, and that's, you know, helpful from a regulatory perspective because that means there are people to regulate.

And, you know, we talk about the code, et cetera. The doesn't fall like manna from Heaven. The code is programmed by people, and, again, these are people we can regulate. And we shouldn't forget that the failure that kicked off the whole series of crypto failures last year was Terra Luna, which was technologically decentralized, albeit very economically centralized in the hands of Do Kwan.

So I just want to sort of offer that as a baseline because we go to a lot of effort in DeFi to get the technological decentralization, and it causes all kinds of problems in terms of scaling problems, et cetera, so there are a lot of challenges. It's basically -- in order to achieve a technological decentralization, you effectively need a more inefficient mechanism than a centralized version because that's the only way decentralization works. So where -- we're picking inefficiency if we're going with underlying decentralized technology, and so we need to think about that in the context of the fact that that technological decentralization is then often overruled by underlying economic centralization. So I just wanted to sort of throw that into the mix as we have this discussion.

MR. REDBORD: Hilary, thank you so much. Steve, I think you were up next.

MR. SUPPAN: Which one? Oh, there we go. I'm sorry. So regarding the issue of technological challenges self-resolving, I assume that includes hardware because you're -- you know, with DeFi, you're going to have a lot more throughput. What's the timeline for resolving the hardware challenges, and, you know, where is -- where are you at in that -- in that process?

MR. ANDRIKOGIANNOPOULOS: Well, I wasn't referring just for hardware. I think there are lots of components in the technology. I was mostly referring to a lot of the protocols and a lot of both software and hardware being early on so that we see the advances in the algorithms and the reliability of the software that becomes more scalable, more reliable, more secure. And at the same time, I think we see the hardware evolving while they're making better use of it, but I wasn't exclusively kind of talking about that. It was mostly, I would say, on the software side because all the algorithms, and the consensus, and the cryptography is basically on the software side of things there, but there are proof-of-work cases where what you're talking about is very relevant.

MR. SUPPAN: Well, just very quickly, you know, the National Science Foundation Grantees Program has done a lot of work on NaN electronics, and, you know, the computer chip that enables that kind of throughput has yet to be invented. You know, a graphene is a very, very unstable element for chips, and yet that is the future, according to the NSF. So that's something I think we need to talk about more because there are some technological limitations that are going to constitute a wall if they are not resolved, and I'm not sure if they can be self-resolved, as it were, through protocols.

MR. REDBORD: Thank you. Thank you so much. Todd, then Dan.

MR. CONKLIN: Thanks, Ari. So at the start of the Russia invasion of Ukraine early last year, there was a lot of talk in in media sources and also questions from the -- from the Congress, in particular, about the use of DeFi to potentially enable sanctions evasion. And Treasury was very clear early on to state that we weren't observing any of that activity, and it wasn't a particularly significant concern given the scale of Russia's typical evasion activities. Is there anything that anyone's observed that should warrant Treasury adjusting that viewpoint?

MR. REDBORD: I'll take a quick crack at it and then kick it over -- kick it over here. I think it's interesting. I think that we were all getting those questions in the wake of the invasion of Ukraine, and I think the consensus answer was, no, there's not enough liquidity in the entire crypto market to run a G20 economy overnight, as you so eloquently said at the time, and still agree with that position. I think what we're seeing is the attempt to use cryptocurrency to evade sanctions at the margins, right, in much smaller amounts by paramilitary groups and, you know, others trying to raise cryptocurrency to support the war effort in much, much smaller ways, using the non-compliant VASPs that Treasury has been going after for the last year or so, but, again, nothing sort of Kremlin or Russia writ large, just sort of much smaller, on-the-margins types of groups. Adam?

MR. ZARAZINSKI: Thank you. So I would echo what you said, Ari, with one exception. So Inca Digital found -- actually, it was -- it was fairly recent, just a few weeks ago -- KuCoin will be providing financial services to sanction Russian banks through their peer-to-peer platform. Volumes varied of course. There was a report that went with it on how Russians use Tether generally, so we're seeing some, but, again, as you said, it's not -- it's not to the degree of like, you know, the entire Russian economy moving to crypto or anything like that. You know, I view it within the realm of everything else that's happening and other avenues for moving money globally, nothing outside of that.

MR. REDBORD: Adam, thank you so much. Dan?

MR. GUIDO: Thanks. Dan Guido, CEO of Trail of Bits. As we're talking about the decentralized nature of these platforms, I feel compelled to point out that at the behest of DARPA last year, Trail of Bits undertook a comprehensive study of the unintended centralities of distributed ledgers and published all our findings in a repeatable manner on the internet. You can find that report. It's, "Are Blockchains Decentralized: Unintended Centralities in Distributed Ledgers."

We discussed a large number of different types of unintended centralities that cover a lot of what Hilary said and more, and it includes empirical data that, again, is reproduceable based on our observed state of many of the most popular blockchains that are available right now. And the findings are, I think, aligned with Hilary's statements that there are significant unintended centralities, that the privileged set of entities that exist in this -- in this industry are numerous, and the opportunity to manipulate the operation of these blockchains by co-opting them is quite high. So I'll let folks search for that report and grab it.

MR. REDBORD: Thank you so much for that. I think one more and looking forward to, like, way more robust conversation as we continue today and over the few months. But obviously, our time is short today, so, I mean, you're going to take us home here for this session.

MR. SUPPAN: Sure. Thank you, Ari. Let me start by thanking you for pointing out and kicking this discussion off by pointing out the fact that the SBF failure, the FTX failure was not a failure of crypto. And thank you, Nikos, for drawing the connection between democratic access, and democratic principles, and openness of blockchains. One main thing that I want to bring up to -- for discussion to everyone is to recognize the fact that we're at the cusp of a technological shift. For many, many, many years we were beholden to single centralized systems. I think some of the people in this room have experienced interacting with mainframes, and from that basis, we went to client server systems. Almost all of the services that we're familiar with -- the Facebooks of the world, the Googles of the world -- are client server systems where we are essentially the serfs and somebody is providing the service to us as centralized entities providing the service to us.

And we're now at the cusp of the emergence of Byzantine fault tolerant systems where there is no server operator, where the service itself is comprised of a bunch of people coming together and holding up that service without coordination, without a single coordinator. That in itself is incredibly empowering and very, very exciting, and that's, I think, the thing that has brought us here.

In coming together like this, one of the main things that we need to be cognizant of is decentralization theater. I think there is true decentralization that lies at the -- at the heart of all this technological change. That's what gives it its power. But at the same time, we need to be -- to be very cognizant of systems that appear decentralized in nature but have these centralized components that are quite concerning.

It is a common belief that the industry is against regulation in this space, and I would like to reiterate that that is not true, that there are many players in the space that would like to see the dream carried out in full, that we believe that these systems gain their strength from actually being truly decentralized, and that the role of regulation here is to ensure that the democratic principles, that the decentralization goal remains upheld. And that, I think, is going to be one of the challenges for us going forward.

There are some falsehoods, there are some myths that are commonly repeated here. One of them is -- has to do with proof of work versus proof of stake. I think that got brought up earlier today. In proof of work, people use U.S. dollars to buy -- to purchase mines from China, then those mining equipment -- that mining equipment then creates new coins. In proof of stake, they use another form of currency, typically Ether, to stake, and then they get new coins. The two processes are exactly identical as long as the cost of entry or the process of entry is open. So there is indeed something to pay attention to, but that thing is not the form of payment. It is whether or not entry into the system is open to all. So with that, I'd like to wrap up.

MR. REDBORD: Thank you so much for those comments, and I will now hand things over to our chair to present the next session.

MS. HOUSE: Thank you, Ari. Really appreciate the discussion here, and, Justin, I did see your flag go up late, so I'll turn it over to you to kick off the next conversation -- sorry -- after our next presentation. We do have two more modules or presentations related to DeFi coming up, one on identity and the next on exploits and vulnerabilities.

But I did want to just highlight some of the really interesting themes that just came up in this discussion that I think are so important, and some of them hearken back to some of the critical points that Commissioner Goldsmith Romero mentioned.

Accountability was a critical one that came up in our presentations there: what does that mean; how "de" is the "fi" in "DeFi;" issues around governance, challenges, and benefits, and promise for open-source software; old finance/new tech; tech being able to solve policy problems; issues around illicit finance; extensive decentralization, et cetera. So all that is a really interesting foundation to guide our continued discussion and this next presentation.

So for our third presentation regarding DeFi issues, I will jointly present with Jill Gunter, Chief Strategy Officer of Espresso Systems, on the topic of Digital Identity, Privacy, Non-Hosted Wallets: What's on the Horizon. So I'm going to kick off our presentation and then turn it over to Jill to continue on, and then I'll help close up, first, by highlighting that I deeply appreciate Commissioner

Johnson referencing NIST. Of course, I love NIST and am very excited about Kevin presenting later today because they've obviously played a really leading role in establishing digital identity guidelines in the U.S., even currently have a new revision of the 863 Series, which I'm sure everyone here knows and is going to comment on the during that period.

But identity sits at the heart of finance, of consumer service provision, and crypto. In a world of public ledgers where not everyone necessarily wants to publish their transactions on public ledgers, and some of you that -- in fact, the transparency, some of the great benefits for investigations that Ari spoke to earlier might actually be more of a bug rather than a feature of the system, and some fascinating technological innovations that are currently happening in the world of privacy-enhancing technologies, including that Jill is helping to lead at her company.

The future of these systems may not inherently be transparent or will have some interesting balances that will have to occur between what information is disclosable and transparent on the ledger, and then how do you ensure things like accountability and discoverability inside of the systems that are obfuscated with privacy-enhancing technologies.

So to kick off that discussion, if we can move to the next slide. Thank you.

First, I'll just start with a general overview of identity. I promise not to try to turn everyone into a digital identity expert today, although I would love it. There's already many in the room and excited about Jill's comments later, but first to highlight that identity is -- it is complex. It is big. It's a concept. It's technologies. It's processes. It's regulations. It really -- the context of you talking about identity also is critical to understand then what is the identity that you care about.

As a taxpayer or for myself as a veteran, when I go into the VA, those are different identities than if I'm going to a gaming conference and someone cares about my gamer tag and my accuracy scores at games. They're terrible. I don't -- I don't compete in any gaming conferences, but that's different than when I'm trying to get a line of credit or if I'm a beneficial owner for -- and I'm registering a company.

Some of the -- some of the terminology that you'll hear in the context of identity includes attributes, so related to my identity, what are some of those features or elements of my identity and who I am. It could be things that we -- that we think of and use for more official identity and how we interact with the government, which could be my name, my social security number, my address. Certain things and features, attributes that the government might actually be more the authoritative owner of that attribute, like my social security number and identifier, or an attribute that has been used prevalently, including in the financial system, to potentially be used as an authenticator. We'll get into that in a second, but also it could be your credit score. It could be your gaming history.

Evidence is the kind of thing that you present to prove that you that identity is real and that it belongs to me. When we talk about KYC, evidence is something that all the financial institutions here are very used to having to consider, whether it's documentary verification or non-documentary verification. What is that evidence that will attest to the fact that my identity is mine and that it is real?

So underneath identity also is this concept of assurance, the confidence or the strength in that identity being real and that it is mine, in fact. So first, you have identity proofing and enrollment, and I -- part of why I wanted to highlight all these things, really just to help set the stage for identity being complex, is also that the ways that identity is exploited implicates different solutions. Identity proofing, when it is exploited, which is basically going -- think in the context of financial services. If I walk in to on board at a bank, synthetic identity fraud is a great example of identity proofing not going very well. And it's me exploiting the fact that an institution might just accept my name and a social that I've purchased for 18 cents on the dark web or made up and hope that it's real, or that there isn't a very strong proofing process on the back end with an address. So that is an example of exploitation of identity verification, but I can strengthen that with more and stronger evidence and higher assurance levels consistent with NIST standards.

Authentication, a different thing. That's when I'm trying to use a credential, like a username and password, potentially weaker authentication or something stronger, for all feds in the room using a PIV card or a CAT to access a system. Having multiple factors can lead to a stronger level of authentication, factors being a variety of three things -- either something that you have, something that you are, or something that you know -- having at least two of them, if not three of them, to strengthen your authentication. And if that's compromised, like for -- if you've -- if your information has been compromised in several breaches and a cybercriminal has purchased -- has purchased or stolen, in fact, those credentials and used it in the conduct of fraud, account takeovers are an instance of that being exploited.

So again, the type of exploitation and the solution to fix it is different if it's authentication versus if it's verification that has a weakness in it and that's being exploited. And then with federation as well where it can be exploited through assertion, modification, or redirection, but you can strengthen that through stronger trust agreements inside of that federated enterprise where the identity is being used or injection protection.

Some considerations that Jill will speak to and some of the interesting solutions in the DeFi space: security. Of course we want stronger assurance and security in our identity systems, privacy -- anonymity and privacy not being the same thing. Typically, "privacy" means there is data that is discoverable or disclosable under certain permissions, protections, and conditions, but what is that information that should be disclosable, and how and to whom should it be disclosed? Usability and equity also a key factor in the -- in the updated guidance that NIST has published for comment.

And then finally KYC. "KYC" is a common term I know that everyone will be familiar with here. It's really more of a -- of a broad term that points to a lot of standards and regulations related to knowing your customer. That points to some of the different elements of identity because it means, you know, establishing the identity and forming a reasonable belief that it belongs to my customer and that it is real, but then also other information related to "identity" being a broader term. Understanding the risk profile of that person, that's different information. That's watching their transaction history, conducting due diligence to understand the broader risk profile.

So now that I've set the stage for what is identity, I'm going to turn it over to Jill to walk us through some of the DeFi identity landscape solutions.

MS. GUNTER: Thank you, Carole, and I am very sorry that I can't be there today in person, but it's a real privilege to be able to present alongside Carole here, nonetheless. So I will start by walking through some of the landscape around identity products as they exist today in Web3, then move to privacy and finally self-custody.

So I'm not sure if you can see the slides up here on the screen, but if we could move to the next slide on the DeFi identity landscape. There we go.

So today, you know, this goes to show that there are many working products that exist out there that are widely used by users, that users are gaining benefits from -- every day within the DeFi world and within Web3 in general. Some of these are still under construction, as it were. Some of these are out there live. So we have identity products that are working on the compliance and KYC front, creating attestations for wallet addresses -- we'll get to wallets in a couple of minutes here -- and enabling a compliance layer to exist, again, within the DeFi and Web3 landscape.

There are projects working on civil resistance, so being able to guarantee that one wallet maps to one human in the real world, being able to guarantee that people are not able to create, you know, multiples of themselves as representations in this digital space. For example, being able to create many wallets that actually, again, just map one person and be able to claim rewards and things like this through that kind of mechanism. Along these similar lines, there are many projects working on universal basic income for which civil resistance is a very important quality. So Worldcoin, Proof of Humanity, these are a couple of the projects working in this direction.

We, of course, have standards bodies, some of which are Web3 specific, others of which, like NIST, of course, you know, creates standards that run much farther and wider than just Web3, but we in this world still reference them heavily. And then finally, there's a whole landscape of projects working on reputation products and protocols that that map to identity as well. As Carole just covered, reputation is but one facet really of identity.

So if we can move to the next slide here, I'll give a very brief rundown of one example identity product. We're going to focus in here for a moment on the Ethereum Name Service just to give the folks in the room a sense of some of the value that users can get out of these types of identity products as they exist today.

So with the Ethereum Name Service -- this is a screenshot actually of my own Ethereum name -- the long string you see across the top here is 0XD0C, et cetera. That is one of my Ethereum wallets, and I have mapped that to an Ethereum name that is human readable: JRG, my initials, dot-eth. And with that, I can self-identify and publicly affiliate myself with an Ethereum address. For a long time I had JRG.eth posted on my personal website, my Twitter profile. People even put it in places, like, on their LinkedIn, and that allows me to prove things about my on-chain activity.

I can show to the world that I donated Ethereum to the fundraiser to support Ukraine last year. I can show to the world NFTs that I've collected or DeFi apps and protocols that I've used, and I can also connect into an increasing landscape of decentralized social applications using my ENS identity. So this is just, again, but one example of products that exist today that, again, users are deriving, perhaps limited right now, but increasingly growing value from.

We'll move on to privacy here and similarly run quickly through the privacy landscape. So there are many different privacy products that exist today. I know that privacy can be a somewhat scary word when it comes to crypto and Web3, partially because, as Ari laid out in the previous session, you know, one of the great benefits actually to regulators and enforcement officers, but also to just users of these systems, is that they are, by default, transparent. You can see all of the transactions taking place. This has been a great boon to the industry in being able to track down and clamp down on bad actors using it for illicit purposes.

However, the transparent nature of these systems also greatly limit their usability and their applicability to a whole host of use cases. It limits their applicability to create value for institutional financial actors who are going to be much more sensitive around data disclosures. It limits the ability of these protocols, and projects, and products to create value within the payments landscape since, generally, having fully-transparent payments is an unacceptable feature of a true payment system being used for anything, ranging from payroll, to cross-border payments, remittances, things of this nature.

So we have, again, this full kind of spectrum or landscape of privacy solutions. There are private payments protocols, including the company that I work for, Espresso Systems, but also including many others, like Zcash, Iron Fish, a recently-announced product called Privacy Pools, that aim to create options for private payments but with a compliance emphasis, so strongly emphasizing compliance tools to go along with the privacy being offered.

There are also, it's worth acknowledging, plenty of products out there that emphasize private payments that take a different approach. They are emphasizing full privacy no matter what the circumstances. And, you know, I think if you spoke to the creators of those products, they would think that, you know, this is a reasonable thing to do sort of in the defense of transactional freedoms. And, you know, for the most part, I think that these folks are focused on the types of users, like dissidents in places like Hong Kong and so forth, but there's obviously a much broader conversation to be had as to whether the tradeoffs being made to enable full privacy are acceptable or not.

There's also a whole landscape of privacy products that then plug into other products within the Web3 ecosystem, so privacy-oriented DeFi enabling traders and users of DeFi applications to be able to mask their positions in order not to get front run. There are configurable privacy products like the one that I'm working on, and I'll get into that further in a moment, and then there's also private smart contract systems. Aleo and Aztec are two examples of these.

So if we can go to the next slide, I will run briefly through an example of how privacy can be a spectrum within, you know, even one given product within Web3. It doesn't need to be fully black or white, you know, all-or-nothing privacy.

So CAPE is a product that I've helped develop. Within the CAPE, which stands for configurable asset privacy, asset creators, for example, stablecoin providers, can create versions of their assets that have customized privacy guarantees to meet their risk requirements. So as an example, a stablecoin provider can go in and use the interface that you're looking at here as well as the contract system that's running it to generate a version of their stablecoin or a wrapper for their existing stablecoin, which is private to the general public. But the stablecoin organization themselves can retain what we call view keys to allow them to have full insight into the full transaction graph of addresses, amounts, and so forth, just as if it was happening on the transparent blockchain. We aspire to unlock use cases like payments, like institutional-friendly DeFi, while still enabling, again, the parties that need it to manage their risk requirements, including those of compliance.

Finally, because this presentation is not just about identity and not just about privacy, we'll move on to the self-custody landscape, also known as unhosted wallets. Again, I think "unhosted wallets" have become kind of a scary word within crypto. Really what we're talking about is wallets that enable users to custody their own assets without reliance on a middleman, without reliance on an exchange or third-party custodian.

So across the landscape here, we have hardware wallets, things like Ledger or Trezor. I know many of you in the room are familiar with these things, but, you know, it's -- basically, it looks like a USB key that enables someone to hold on to their own bitcoin, Ethereum, or other tokens. We also have browser-based wallets that pop up as a sort of Chrome extension, like MetaMask, Coinbase Wallet, WalletConnect. And then we, of course, also have mobile wallets, so things that -- applications that can run on your iPhone or Android.

And then finally, on the more institutional side, some choose to use multi-party computation wallets, which effectively means that it's not just a single party who is custodying the private keys that gives themselves access to their assets, but that key is actually split between a host of users that have to come together in order to gain access to those assets. So again, you know, I know many of you in the room will be very familiar with MetaMask, but I'll just briefly run through an example of such a wallet product, partly to demonstrate that, at least as a user, it doesn't feel like a scary thing when it's sitting in front of you.

You know, MetaMask is one of the most used self-custody wallets out there. It offers a gateway not only to assets but also to decentralized applications right there as a pop-up, plug-in that comes up on your screen. When you open up a DeFi application or, in this case, what I've shown a screenshot of is a social media application called Mirror.xyz. It's very much like medium or any other blog-hosting platform, but my gateway into it is by logging in with my MetaMask account, which, as you can see, has popped up on the screen, and I sign in, and then my wallet address or my ENS name, JRG.eth, then can be associated with all of my entries in the blog. And so the goal here was really just to walk through a handful of examples of these products and to give the committee here some familiarization with the landscape.

We'll go to the next slide here.

And just again, you know, these three topics -- identity, privacy, and unhosted wallets -- have been heavily in the news with many issues around them. But

I would encourage, and my hope for this conversation and the conversations to come over the course of this year and beyond, can be around how we can foster innovation on all three of these fronts but do so in a way that meets our regulatory and policy goals. And with that, I'll hand it back over to Carole to carry us out. Thank you.

MS. HOUSE: Thank you, Jill. Really appreciate that. So the final two slides I really think will help kick off some of the discussion, and, Justin, looking at you to kick off the first comment. I'm sure you have some in response to identity. But first, some of the areas that Jill and I, as we were thinking and brainstorming on some of the areas that would be beneficial for examination on the technology and policy side to provide standards, to drive more clarity in this space, some of the things that we thought to finding the key features of what an identity system should be looking like in the DeFi space. What do we want to see? What do we not want to see? Issues related to portability, verifiability, equity of access, privacy, appropriate privacy, recoverability. If your identity gets stolen, like, what is the recourse for victims, which is too often overlooked in certain DeFi systems, establishing what the right use cases are that we care about, looking at traditional identity fixes needed for DeFi.

Part of why I started off with that, I'm sure riveting to the discussion about traditional identity and what it looks like outside DeFi context, is that there's a lot of issues in the identity space outside of just DeFi. And while there's wonderful innovations that are currently going on in the technological space related to decentralized identity -- interesting standards, all the products that Jill spoke to earlier -- some of those issues or most of those innovations are not, or at least some of them are not necessarily looking back to the current problems that exist in the traditional identity space and that may inherently just end up ensuring that we're then importing all those problems from the traditional identity space, and then further decentralizing it in an ecosystem where accountability is higher of question.

So what are the kinds of issues that we can -- so that we can make sure that synthetic identity fraud, which is rampant in the banking system is not, in fact, rampant in the DeFi ecosystem? What are -- what are those kinds of solutions that could be under way, that could be put under way on the government side versus on the industry side and market developments? How do we ensure responsibility in the ecosystem?

What does responsibility and accountability -- again, pointing back to Commissioner Goldsmith Romero's comments in her opening remarks, what does accountability properly look like in a decentralized identity ecosystem? Can you actually have that without regulating the providers of those -- of those identities if you want to ensure that no other stakeholders inside of the decentralized finance system are regulated, but then, ultimately, when victims are hurt or there's a national security threat, authorities have to go somewhere. Is that a role that future trusted identity providers should fall underneath in the regulatory landscape?

And finally, on the privacy side, how do we incentivize development for data protection with developers to ensure appropriate discoverability? How do we encourage our builders? What are the right incentives, both sticks and carrots, to protect user privacy without sacrificing the ability of government and other appropriate authorities to get access to that critical information? How do we prioritize and promote tech with protections without condoning products that -- threatening actors. And then for unhosted wallets, specifically, we need to calibrate the role, treatment, and freedom, and responsibility of builders. How do we avoid undue burden and what is undue burden on the developers for open-source wallet?

And then finally, in examining risk, accountability, and discoverability in evolving systems in a world of unhosted wallets, the final question was, what is the right kind of identity system that allows for that proper discoverability of certain information and to which counterparties and authorities for unhosted wallets in a system that currently relies largely on central counterparties and cash-out points, as Ari spoke to, that's currently the landscape? However, it's an assumption that that will be the landscape forever, right?

Part of the vision of DeFi is to do away with the need for cash-out points and that it becomes a self-sustaining ecosystem, or vendors will accept these assets in exchange for goods and services. So when you can no -- when you can no longer rely on those central parties and cash-outs, that moves on to be what does the right decentralized identity ecosystem look like in a world that can't rely on those centralized parties? I'm going to have less need for those cash-out points.

And then finally, another really interesting aspect of these ecosystems is the fact that you have both financial assets and non-financial assets riding the same rails in an interesting world where, like, in our worlds of information transfer and internet activity, we don't make identity an inherent part of that activity, you're able to establish trust and identity across the internet. But it's not required and always tagged onto your activity versus in the financial rails, where on traditional rails, identity is always there. It's a part of how you manage account services, et cetera. But now, the same rails will support information transfer and value transfer, and with the right obfuscation you may not be able to tell which is, in fact, occurring. How do we make sure that proper -- that proper identity information is available for financial information but you preserve privacy for non-financial activity?

So that closes out our discussion -- our discussion and our presentation on identity. Now, I would love to open it up to the floor for the different TAC members for questions, comments, and reactions. Dan, I think you were first. I'd love to turn it over to you.

MR. AWREY: Thank you so much, Carole, and to the commissioners. This has been a great event so far. I have a factual question, I think, for Jill and an explanation for why I think it's an important question, and it had to do with the compliance layer components. And I'm wondering to what extent existing strategies for developing this compliance layer rely on centralized on and offramps. So the example that was given, the CAPE example, easy enough to understand in the context of a centralized stablecoin that is subject to legal KYC AML obligations and then wants to use products that enable it to comply with those obligations.

But the direction of travel, if the earlier discussion is to be believed, is that, ultimately, you know, as a practical application of Metcalfe's law, as more and more people use unhosted wallets, the on- and off-ramps are going to become less important. And, in fact, the on- and off-ramps themselves may become decentralized over time, which then means the question of reliance on on- and off-ramps for the compliance layer becomes very important to design up front, knowing that you may very well be doubling down on an existing compliance strategy that's not fit for purpose in a decentralized network.

And then I just wanted to flag how big a sea change that is from an AML KYC sanctions perspective, right? Risk-based AML KYC laws, basically, okay, there's a centralized actor over here. We're going to come down on you hard if you don't manage these risks, so manage them. It seems to me that that challenge and who you're placing the burden on then becomes fundamentally different as you move towards a system of decentralized actors. So thank you.

MS. HOUSE: Great question. Jill, reaction.

MS. GUNTER: Yeah, thank you so much, Dan. That's an excellent question, and I appreciate the forward-looking nature of it as well because I think that this is going to become more and more of an issue as more and more value is just transacted solely on chain. And increasingly, you know, if these products do continue to take off, we will see less and less need for actors to be cashing in and out, and there's, of course, fears that we're already seeing that in some cases.

I'm going to point to Circle, and I know that we have some representatives of Circle in the room, I think is a great example of a product that's been created with this exact feature in mind. And so we can look at examples where hacks have occurred, where the exploiters have been able to extract USDC from vulnerable contracts, and Circle is able to move in swiftly and halt the USDC from any further movements.

They're able to freeze it on chain, and that's an example where that can be enforced on chain, again, without having to push it to the on- and off-ramp itself.

There again, of course, you're still relying on a centralized actor who's taking on that responsibility. If you look at the way that these hacks tend to go down, one of the first things that the exploiters tend to do if they're savvy is to try and trade out of their USDC, knowing that Circle is going to be this kind of responsible actor that does engage in freezing the assets, and they try to move into decentralized assets where there is no such actor.

That, I would highlight, is not strictly a privacy problem. That is a problem that exists on chain in general. Again, we can at least trace it outside of the privacy context, but this is why we, at least within my company, have emphasized our product to build around these types of centralized actors. And I think that there is a contingent within the crypto community that believes that those centralized actors are going to continue to grow in their importance and influence because if you look at what mainstream users want as well, and this is, I think, an important consideration, not only in the product that I've developed but also in products that emphasize compliance in general.

Mainstream users do not want their funds being mixed with North Korea. They don't want to be aiding and abetting illicit actors. They just want to be granted the sort of baseline privacy that they have when they are, you know, using the traditional banking system. And so that is where a lot of my optimism lies around this being solved is in actually the demands of users. But it's a great question. I welcome it, and I look forward to further conversation on it because I don't think that we have the full answer as of yet today.

MS. HOUSE: Thanks, Jill. And, Justin, I said it twice and then missed you. I'm sorry. I would love to turn to you for your reaction.

MR. SLAUGHTER: Happy to take it. I think Corey was happy to take it. I think Corey gets first step because Circle was mentioned, unless you don't need to say anything. All right then. Thanks so much for this. Of course it's great to be here with all these leading luminaries.

I just want to make two points. First off, Carole, you have a wonderful litany of questions on what we can discuss in this upcoming series of meetings. I can't cover them all, but I will say I think the most important thing to focus on is the need to do things like this to engage between industry stakeholders, nonprofits, academics about the technology and its base rather than simply wait for it to either go away or to develop on its own. Without that kind of active engagement, it's likely that choices will be made for everybody by default or by accident.

On the subject of digital identity, I also wanted to stress one thing that I think gets often overlooked. Too often people suggest that only people overseas -- activists in Iran, people fighting for freedom in Ukraine -- need privacy, especially for activism. I think that is an incredible red herring. In this country, and a lot of us know this, the most

-- one of the most dangerous jobs you can do is activism. If you're a union organizer, your privacy is critical because there's a real risk your ability to organize your union. Your workplace will be broken the moment you work with somebody else and they know you. That is, in fact, I think, the most interesting digital identity startup I've seen so far, which is these DAOs, like democraDAO or work by the Blockchain Social. It's focusing on how you can build organizing power through blockchain, through DLT, through crypto in a way that increases the power of workers versus the very powerful people at the top.

The other thing I was going to note is to respond to the first panel. I think it's really good to see that we all agree the decentralization generally is positive but that there's a lot of fake actors in this space. That is, in fact, probably one of the best role for regulators. Encouraging the industry to move toward decentralization is positive. That's one of the ways you can channel the growth of this industry because I do think if left to our own devices, we could become a replication of a lot of traditional finance.

People often forget this. There's only three or four companies that run almost every major financial market in this country, and they're great companies: CME Group, Intercontinental Exchange. There's been a substantial lack of competition there, and one of the reasons I find this space so interesting is the chance to do this again hopefully and encourage more competition, encourage more growth. But that's something that can only be done with the hands of the regulators.

MS. HOUSE: Thanks so much, Justin, and I did want to reinforce the point. I know Tony mentioned earlier if you'll identify yourself and where you're from. Thank you. Justin Slaughter from Paradigm. So next, Dan. Dan, if you can introduce yourself, yeah.

MR. GUIDO: I'm Dan Guido from Trail of Bits. I wanted to make a point to just highlight the extreme level of technical challenge that folks like Espresso Systems and other people doing work in the zero knowledge and privacy-preserving cryptography space are underneath. This field is completely unstandardized, yet it underpins a lot of the key features that are required to make things, like digital identity, and what else the folks from Paradigm just described. These are -- you know, there aren't verifiably good implementations of them, and a lot of times when people are trying to build them, they have to go back to the original papers as they were written by academics in the 80s, 70s, 60s, who weren't aware of how we would be trying to use them today.

So we found a number of vulnerabilities in these systems at Trail of Bits, and I've had to report them to others, where people have actually just straight-up followed what are in these academic papers to the letter, but some of the descriptions of them in the -- in the papers themselves have been broken. So there's really, like, an extraordinary amount of technical expertise required to build privacy-preserving encryption systems that are defensible, that are reliable, that are resilient against attacks.

We've made a small contribution here in the absence of that standardization. Trail of Bits published a resource called ZKDocs.com that prescribes known good solutions for actually implementing these technologies to make it easier for people, but just really wanted to highlight that these are -- these are research systems that now we have put into public practice. And it feels like a lot of people are sort of dancing on the lip of a volcano when it comes to using them to safeguard the privacy of others.

MS. HOUSE: Thank you, Dan, and I appreciate the insights and look forward to your presentation in a moment. Michael Shaulov?

MR. SHAULOV: Thank you. So, first of all, thanks for a great overview of the colleagues. As you guys have probably seen, I'm Michael Shaulov. I'm the CEO of Fireblocks, and we actually play across multiple of those layers in the stack, so I have quite a few things to kind of highlight.

The first one was around unchained identity, and I think that this is probably the most critical part that we need to sort out and create some level of either a sandbox, or guidance, or regulation that will be -- where a compliance officer will be -- we start to be comfortable with this, right, because when we think about the promise of DeFi and the promise of the democratization of finance, access to all those protocols in a compliant way, right, the most important thing is that as assets are being tokenized or assets are being blocked -- being brought to the blockchain, we want to introduce consumer protections, right, and investor eligibility. And the only native way to do it, actually, at the large consumer base is through on-chain identity.

What we've experimented and what we've seen so far, and, you know, we've done one of the initial projects in the space with a partner called AVA, what's called Aave Arc, is that when you approach compliance officers across regulated firms, whether those are banks or asset managers, right, the -- although the concept of unchained identity sort of somewhat technically understood the fact that there is no regulatory framework or at least sandbox that they can get the guidance from, is sort of deteriorating those compliance officers from starting to explore and give sort of a green light for the asset managers, or the banks, or the other financial institutions to start experimenting. And, therefore, I think it's critical to create some level of, you know, regulatory clarity around how people can engage because, otherwise, at least, like, the institutional space is somewhat stuck.

The second aspect that I did -- I also wanted to mention is that, probably similar things hold for the hosted wallets, especially as it comes to the institutional space. So, you know, we are -- we are a player in the, we call it direct custody, but, effectively, it is a noncustodial world for institutions. We have a very large scale of users, about 1,800 institutions, that are using it at the moment. And interestingly enough, I think that when we look at the counterparty risk, right, of a custodial service vis-a-vis noncustodial service, and especially after, you know, what we've seen the last couple of weeks with Silicon Valley Bank, the appetite and the view of many large asset managers, and, you know, other institutions, and fintech players is that noncustodial wallets, unhosted wallets, direct custody has to be well defined. And it is an opportunity to reduce the counterparty risk from what currently exists with the -- with the -- in a traditional financial market.

You would be surprised how many really large asset managers call us and sort of ask us if we have a view on how in the future they can do self-custody for, you know, government Treasuries or other assets, right, which currently can only be held through large-scale custodial banks or centralized depositories. So I think that those three issues has to sort of be front and center in order for us to make advancements in the space. Thank you.

MS. HOUSE: Thank you, Michael. Appreciate it, and I knew there would be a robust discussion. So I'll close off additional flags being raised, but the order will be Michael, Emin, Michael, Sunil, and Ben. So, Michael Greenwald, if you'll kick us off.

MR. GREENWALD: Thank you. Thank you, Carole, and Commissioner, and Ari, for having us here. Todd, you mentioned Ukraine, and I was just curious, from a digital identity perspective, what lessons learned or best practices have we seen given the growth of humanitarian payments with digital assets and digital identity. Is there anything we've captured or learned from those experiences?

MS. HOUSE: John, I don't know if you have any reactions. I know that there's a lot of use cases certainly at least being claimed from some folks in being able to provide access to financial services to lots of disenfranchised individuals. In fact, in the context of Ukraine, it was even pointed to by, I think it was the Secretary of State or it might've been DepSec that made a comment after the issuance of the executive order on the benefit of humanitarian relief, one of the goods that we see. She also highlighted some contrasts and some negative consequences that we've seen, of course, using crypto.

But I also know that a key point that I've seen raised, certainly by those who are conducting those humanitarian services, highlight to those who just don't have access to -- regular access to identity documents and credentials, and basically being able to provide them access to financial services in some of those -- in some of those areas, even though they don't have some of that traditional identification.

MS. GUNTER: I'll just --

MS. HOUSE: Jill, if you have anything?

MS. GUNTER: I'll just -- yeah, I'll just chime in as well. This is not on the identity side, but I think it's worth noting that one of the big narratives that grew out of the humanitarian aid donations to Ukraine was actually the importance of the tool, Tornado Cash, in enabling lots of people to feel comfortable making such donations. Vitalik Buterin, who is one of the co-founders of Ethereum itself, came out and said himself that given his background and nationality, he was primarily comfortable at the time donating to the Ukraine efforts on chain with the Ethereum because he could use Tornado Cash as a privacy tool.

Now, Tornado Cash since, of course, has been shut down by OFAC, sanctioned by OFAC due to its role in helping to wash an outsized amount of assets that have been hacked, it's believed by the Lazarus Group and North Korea, and so that, of course, is the problematic side of that product. But I do think, since the question came up around what types of products were used around the -- specifically the use case of getting aid to Ukraine, that was one product that was very much highlighted throughout the discussions around what enabled that to happen.

MS. HOUSE: Thanks, Jill. I did see Corey's flag go up, and I suspect it's because Circle has an example to give.

MR. THEN: Sure, yeah. Thanks for your observation, Michael, and for having us here. I just wanted to point out, like, this is one of the really exciting things in uses for crypto. Circle has a partnership with the United Nations High Commissioner for Refugees right now where we're piloting sending USDC over the blockchain to displaced people, right? And so this was an extreme, like, vetting exercise with the U.N. to, you know, speak to the strength of this technology.

And in addition to, you know, benefitting end users who can either take it and receive USDC on their -- on their device, perhaps sitting in a basement when there's bombs going off outside, and then carry that throughout the country, there are on- and off-ramps, cash in/cash out, that's required if you want physical money, right? But a lot of that money, we believe, will just circulate on chain, and we think that that's going to increase over time. We've seen it with a separate partnership that we had in Venezuela to get USDC to frontline workers who weren't being paid by the Maduro government, and partnered with Treasury and State on that.

The other advantage outside of just to an end user, if you think about an aid organization like UNHCR, is you can see precisely where the funds are going and make sure that they're going to the right people that have been vetted to receive them. That stands in stark contrast to, you know, I had a friend who is in the U.S. military who, during Iraq reconstruction, was literally putting his own life at danger to drop off pallets of cash, a very high percentage of which walked away into the wrong hands, right? So this is a really, really promising use case that, I think, is unappreciated by people outside of the community but that we're going to hear a lot more about moving forward.

MS. HOUSE: Thank you, Corey. Appreciate that. In the interest of time, so Emin, Sunil, and Ben, and if you'd rather reserve your comments for the -- for the next discussion time after the exploits talk, just let me know. But Emin?

MR. SIRER: So very quickly, I'd like to make three technological observations. I think the very first one is the observation that the name service that is so successful and is behind the success of the internet is a federated one. And I posit to you that any kind of a centralized attempt or any attempt to centralize the naming on the internet would have ensured the failure of this thing that we take for granted that powers much of our economy and the global economy. So that's one, and it's a federated system. It's not a fully decentralized system. The name service, DNS, is a federated one, but it's not fully decentralized in the way that blockchains can be. The technology back then was different, and the best we could have done back then was federation and not full decentralization.

Second observation is that the rules required for compliance, for identity-related compliance, vary greatly from state to state, let alone from country to country, from economic bloc to economic bloc. So it's a -- it's a -- it's a pipe dream to imagine that there could be one set of rules that could apply across the globe. So any attempt to try to regulate in that direction towards a single set of global rules will probably be mired in so many meetings and so little outcome that it will probably end up failing by itself. So we have to be able to come up with techniques and technologies that can be applied separately, that can be deployed incrementally, and that allows some autonomy to various different actors.

And the third and final observation that I have that I'd like to posit to the board here is the evolution of the chains that we take for granted. I heard a lot of discussion about chain, on chain, singular, but if I look around at what's happening, we did start out with networks like bitcoin, which is a single-asset, single-chain system. We evolved into networks like Ethereum, which is a multi-asset, single-chain system. But now we're in the age of multi-chain systems -- multi-asset, multi-chain systems -- Avalanche, Polkadot, Cosmos. These are recent generation systems that have multiple chains underneath them, and those chains have the capability to have different rule sets apply. We are not hamstrung by the fact that there is a single chain and, therefore, a single set of rules that need to be enforced around the globe. We don't necessarily need to regulate for such a universe, and that, I think, has been an immense source of freedom and ability.

So at our labs, for example, we're building what we call a subnet, an institutional subnet for Wall Street, some of the main players on Wall Street where the compliance requirements that are specific to their needs are enforced on a specific chain for their use. This does have some disadvantages. It means that liquidity is divided a little bit, but it does also have the ability to accommodate regional, specific jurisdictional compliance measures. Thank you.

MS. HOUSE: Thank you, Gun. Appreciate it.

Sunil?

MR. CUTINHO: I have one observation and perhaps a question. So the first one is on DeFi, and I was looking to learn a lot. One of the problems I have when listening to the presentation is that we start with the answer to the question "what," and we spend a lot of time explaining what it is and how it works rather than explaining why do we need it, and in some cases, we lose credibility. I'm a technologist as well. I understand how it works, but when we start saying that it is a solution to a bank run, you lose all credibility. So SVB it was not a problem of transparency. It's a problem of a bank run, a classic, good old bank run. It's just surprising in its ferocity. It just took place in a very, very short amount of time.

The second problem is something that was mentioned before. When I see what's going on in the world on DeFi, I see a lot of centralization, so it loses its purpose, so completely antithetical to the idea of DeFi. So it doesn't make any sense to me, so I'm very skeptical. So a good thing would be for you to convince me why DeFi is important as a solution for financial markets.

Finally, derivatives. Some folks are mentioning DeFi in broad senses, but in derivatives, there is a problem of lifespan of exposure. So if two counterparties have an exposure for an instant, then decentralized works perfectly well. But if two counterparties have an exposure for a lifespan greater than an instant, even for a day, for two days, for a week, for a year, for 30 years, for 50 years, you wouldn't want to face that counterparty. They may be gone. They may actually take the money. They may perform to you and leave to Mexico, or they may leave to another country where there is no extradition laws, so DeFi doesn't help you there. You need some entity to guarantee performance, and that entity enforces contract guarantees. So I think we need to go in pragmatically with DeFi and explain the problems it's solving before we get into what it is and how it solves.

And finally, on identity, as an individual, the thing that is most important to me is control, okay?

How is my identity used, where it is used, and when it is used? And one of the most important things for a solution to provide is that, you know, it needs to be a standard, so it cannot have any frictions associated with it. The moment I have to use -- I have to centralize with one provider, I lose control, so that -- and therein lies the problem with identity, but I do see a great use case for identity. There is a great promise, but I think we need to solve for control. I think as an individual I want control, so that's the problem to solve.

MS. HOUSE: Fantastic insights and also really speaks to, I think, why by Commissioner Goldsmith Romero and the Commission specifically brought together this group because everyone's views are not going to be the same or share the same perspective. So I look forward to that debate and really appreciate that, again, it harkens back to accountability, what that looks like, who are the right intermediaries, or does accountability sit with end consumers. A lot of tough issues and friction points that go on consumers that, honestly, we don't like to put on them either as regulators or as businesses that want to make money by lowering friction and costs for them. So I don't have an answer for it, but, Ben, if in your closing remarks, to close out the discussion, if you want to react to his or her give your own reaction to the identity discussion.

MR. MILNE: Yeah, I think an observation and then just a hope. Just from an observation perspective, I share a similar perspective in that I think we've covered a lot of ground. And it's important to maybe go back sometimes to first principles of just what are the definitions of these words as they relate to market infrastructures or markets generally speaking, because when we say "self-hosted wallets," it means something very different on the consumer side than it means to, let's say, a systemically-important entity.

And so my hope is that there is an opportunity throughout this process to come up with common definitions that we can use going forward, and we can focus on maybe the common definition, starting with what is the traditional finance definition of the word, and then as a DeFi community, trying to think about how those technologies map to the traditional definitions.

MS. HOUSE: Thank you so much. Then for our fourth presentation regarding DeFi issues, we have Dan Guido, founder and CEO of Trail of Bits, and Michael Shaulov, founder and CEO of Fireblocks, will present on the topic of Exploits and Continuing Vulnerabilities in Crypto Markets. Dan will go first followed by Michael.

MR. GUIDO: Okay. Hey, everyone. I'm Dan Guido, the CEO of Trail of Bits. I want to briefly introduce ourselves. So I founded Trail of Bits 10 years ago with the aim to solve the hardest problems in software security. We really believe that we need better tools in order to overcome the challenges at hand, that just best practices are not enough. Over the last 10 years, we've grown the team to about 140 research engineers that are solely focused on these emerging technologies of which blockchain is one.

We work with people across national security, DOD, DARPA, the tech industry. We've worked with companies like Microsoft, Google, Zoom, Epic Games, and we've worked with pretty much half the block chain industry. We have unprecedented visibility into the internal operations and the production of code that occurs in the blockchain industry. And just -- I have to give credit where credit's due. I had my education paid for by the National Science Foundation through their Cyber Corps Program, and very happy to be continuing in public service here.

So next slide.

So one thing that I wanted to make clear is that there is obviously a perception in the industry that no one can get it right, that everything gets hacked every single day, there's two or three companies that are completely obliterated by hacks in the blockchain industry and it's completely overrun with scams, and that security must be an afterthought, these must be the worst possible things that have ever been created and that's why they're getting hacked.

But in reality, I see something different in the everyday practice of Trail of Bits. The blockchain clients that we have are actually the most rapid at incorporating the techniques and the guidance that we give them. They demand it from us. We will beat them up, and they will ask us to hit them harder. It is, however, very difficult for them to understand what they should do. There's an obvious dearth of security expertise across the entire technology industry, but it hits the blockchain industry even harder because the foundations of the field change every day. The kinds of problems that we solved, the kinds of technology that we built in blockchain is different than it was six months ago, then it was a year ago, then it was two years ago. And in order to secure yourself, you need to be perfectly up to date.

And then finally, there's not a lot of information that you can trust. A lot of firms themselves as well as fans of them spread a lot of information that is more marketing and aspiration than it is empirical truth. So that's part of what I'd like to bring to this conversation today is some empirical truth about what we see on the ground, having performed hundreds if not thousands of security audits of these firms.

So next slide.

So the first thing to recognize that's really unique about this field is that it moves far, far faster than -- on technological underpinning than any other field of software. This can make it extremely difficult for standards and practices to apply because if you're using six-months-ago standards, it is completely insufficient to protect you today.

So I've got some examples up here. Before 2020, we actually released on authoritative work of trends that we observed for all the audits that we -- that we looked at. It was 246 findings and tried to determine what are the issues that people have trouble with. That study was conducted, again, in the last year by a different research lab, and the kinds of security issues that affect firms have completely changed. DeFi is now present. The introduction of flash loans changed the risk calculus for a lot of these firms. We're looking at Oracle manipulation, composability bugs. These things simply did not exist a year or two ago.

So this has implications for the way that we compose standards and guidance for these firms, things like NIST CSF, SOC 2, PCI, they're all very high level, and the ones that are very low level obviously going to become outdated immediately. The ones that are very high level aren't going to specifically address the kinds of flaws that these firms need to protect themselves against.

Next slide.

Now, another unique dynamic of this field is that information is public and platforms are shared. In the regular software industry and in D.C., particularly, we've had lots of conversation about building an NTSB for software, for cyber, for technology, and that already exists here. All of these transactions that occur in the blockchain, all of the hacks, they are extraordinarily public, and it's usually your users and other outside firms that find out about them before you do. So this is -- this inverted view on what secret, it's really -- it influences a need for perfection from these firms. You have to take very careful steps forward in order to assure the safety of what you've built because everybody is watching.

So on the right here, we have one of these sort of NTSBs for safety, direct leader board, which is a public collection of memorialized firms that have either been completely knocked out or partially knocked out based on hacks. And many of them have a word next to them, "unaudited," which means that they didn't actually seek outside guidance for whatever changes they made before they made them. So, again, highlighting the fact that an extraordinary amount of expertise is required in order to build these systems safely, and the firms that don't know that and don't live by that end up getting wrecked.

I guess one other really interesting note here is that in the -- in the latest Biden executive order on cybersecurity, we talked about cyber liability, that firms should be liable for the sorts of issues that they produce for the world. That's already here in blockchain, too. These hacks have a direct impact on the finances of these firms, on their governance tokens, on their Treasuries that are rated, so, in fact, they are more motivated. They have not externalized the costs. The costs are internalized for these failures, which kind of drives the reason why these people are so rapidly consuming guidance from firms like mine in contrast to a lot of other clients that we might have where the dynamic is different.

Okay. Next slide.

So finally -- finally -- all of this really influences the need for perfection. Where a lot of other industries can get by with risk mitigation, this field needs risk elimination. You cannot ignore vulnerabilities that were given to you that are simply low severity, right? Everything is high severity, and this also creates issues with these standards yet again. You wouldn't ask NASA to build the software and its rockets that are going up to space with only the NIST CSF, right?

That is what the block needs. The blockchain needs software that is built to precise specifications that always operates the right way. That level of correctness is really not commonly achieved. It really is never achieved by anybody else building software right now. This is what we call high-assurance software. This is the kind of thing you need to do when you're building rockets, when you're launching things into space, when you're building cryptographic libraries or working with software that mediates the life and death of a person.

So I see that, you know, clearly there's a there's a dearth of expertise in this field. There's not enough security experts to get the job done. A lot of people might look at AI as a potential solution to this issue, but that is not the right thing to use. AI, as I've said here, is a paintbrush, and we need a scalpel. We need things that are -- that are precise and algorithmic, not probabilistic.

So really, all of this points to the field needs more research, needs more work, more innovation done to figure out how to secure these systems because they've chosen to try to scale the tallest mountain we have in software security and to do it all right now. After decades' worth of study in computer science, we still don't have a lot of methods, and techniques, and tools available to us to properly meet the bar that we have set. However, there are things we can do right now that we know have to be done in order to get there.

So next slide.

So yeah, just to reiterate some key issues here, that blockchain companies are actually motivated to fix security issues, and they are some of the most security-conscious organizations we have ever worked with. The underlying foundation of the field changes rapidly, and the technology solutions and guidance that you give people six months ago sometimes doesn't apply today. The public nature of these chains and of the hacks that occur is an extraordinary opportunity to learn both for us but also for attackers. So there's extraordinary systemic risk when new attack methods are identified or when new risks are identified that they become exploited by actual attackers within hours, and it can affect the entire ecosystem, and that, you know, we really need to improve on research and innovation here, that best practices are necessary but not sufficient to solve the problem that we've got.

So with that in mind, we do have several things --

Next slide.

-- that I would like to highlight, are extraordinarily important for everyone to meet. This is kind of what we use as one of our internal standards of are you doing the right things to keep your blockchain protocol safe. We developed this in partnership with a number of other firms at a recent conference that, I believe, some people the room may have been at. But these are 12 critical security controls that I think a -- an amateur, or an outsider, or an interested party, or a regulator, a venture investor, a user can have a conversation with -- a productive conversation with a blockchain protocol to determine if they are doing the right things to keep their data safe.

This is -- again, it's not comprehensive. If you could answer all these questions, it still is not enough. But, you know, having a written and tested incident response plan, there's not a world in which that's not a necessity.

So next slide.

I'd like to offer up a couple resources that our company has produced to help further this conversation and add to the safety and security of these systems. That Rekt test has a block of its own. I mentioned a paper that we've described about blockchain decentralization, the empirical data from our audits, as well as the huge number of open-source tools and best practices that we've been able to put out there to help companies do this job well, in addition to, which I forgot to mention, a new AI safety team where we're repeating the same process of building out the security foundations of a field as it's emerging.

So with that, I'd like to pass it over to Michael, who knows a lot about this topic as well.

MR. SHAULOV: Thanks, Dan. Appreciate it, and maybe we can switch to my presentation. Okay. So thanks -- again, thanks so much for having me here. What I want to do in the next 10 minutes or so is to essentially kind of build up from where then where Dan walked us through and give three practical examples of actually, you know, real-world hacks that happened in the last 18 months. And the reason is that because we're sort of like, you know, kind of diving deep into three core issues that I think cover, I wouldn't probably say 100 percent, but, you know, 90 percent of the issues that we see across the -- all the hacks that are on the -- on the website that you guys saw before.

So quick introduction about us. We provide secure infrastructure for financial institutions and Web3 companies in the space. We have a pretty large client base of 1,800 clients that we are servicing. And before actually starting Fireblocks, I spent two decades of my career in cybersecurity, and the last thing that actually led us to starting Fireblocks was that we investigated a breach that happened in South Korea across multiple exchanges back in 2017. So I think that actually looking into some of those hacks gives a pretty good insight of what's going on and what we need to be focusing on.

Next slide.

So the three facts that I want to kind of review is -- the first one is basically the hack that Haplin was running. It was brought up earlier by Ari in his presentation, and that's a hack that is related to key management and private key security. The second one is that -- with BadgerDAO, which is a man-in-the-middle attack and related to transaction security. And the last one is something we discussed quite extensively so far is with a protocol called Euler Finance, where this was a smart contract hack. So let's start from the first one, and I'll provide a bit of a setup, and explain what that company was doing, and then essentially what failed.

So Ronin is actually a blockchain that is underpinning a game that was the most popular NFT game called the Axie Infinity. It was operated by a company called Sky Mavis. And in order for you to play that game, you had to basically transfer some, you know, some value into the Ronin blockchain. The Ronin blockchain was effectively four core, basically a copy of the Ethereum blockchain that was used for more high-performance capabilities across that game. And that bridge is -- the bridge was operated in what allowed users, allowed consumers to basically take their Ethereum coins or take USDC and essentially replicate on the running blockchain.

So the way that it was actually operating, without getting into the real deep technical details, is that there was an address that was basically an account on the Ethereum blockchain that if you wanted to play the game as an -- as a user that had some Ethereum coins, you could basically put into that account, and then the Ronin Bridge was basically mirroring those assets on the other blockchain. The way that the bridge was operated, basically Sky Mavis wanted to make sure that it's somewhat, I guess, like a federated approach, not exactly decentralized, but it's not decentralized -- a centralized issue. There are a bunch of validators which are effectively just servers, and those servers, each one of them has a private key that is required to sign transactions for deposit and withdrawal of those coins.

And the way that the structure worked is that there were nine different validators. Five of them were operated by Sky Mavis, another four -- another five by different actors, and you had to have a signature of five of them out of the 9 to basically withdraw assets.

Now, if we go to the next slide, what happened to -- specifically to Sky Mavis is actually a very typical spear phishing attack which, you know, we've seen in the cybersecurity industry for the last, I guess, like, decade and a half. And I think the attribution for this attack goes to Lazarus Group that was mentioned earlier. They're affiliated with North Korea. And in March -- and basically on March 23rd, they were able to withdraw $650 million worth of cryptocurrency from that bridge.

The way that the attack unfolded or the forensics was that what the hackers were able to do is to convince an IT engineer, a dev app engineer, that worked for Sky Mavis to go through a fake interview process for a different blockchain company. Through that interview process, he actually received an assignment and they basically sent them -- sent him a malicious PDF file that he downloaded his computer. That malicious PDF file contained the malware. Once they were basically on the computer of the IT professional, they were able to traverse across all the different servers where they had the private keys for the bridge. And because of some mis-configuration that happened a few months earlier, they were able to also traverse into another validator and effectively collect five private keys, right, that were sufficient to control the bridge and to withdraw the $650 million worth of funds.

Something that I would actually mention about this specific attack is that we see a lot of those examples around private key management where client -- where companies are essentially creating bespoke systems to do that, and they come up with their own practices of how to do key management, although today we have probably, like, you know, a large set of institutional and non-institutional providers that have been doing it for a few -- for a few years already, and underpins a good chunk of the hacks.

Now, let's go to the -- to the next slide to basically discuss the Badger.

So Badger is basically a decentralized application, that what it does, it basically allows users to generate yield on bitcoin using variety of decentralized strategies. And actually the decentralized app is working just fine, but in order for you to access a decentralized application, the most convenient way to do it is through a web application that is provided in kind of Web 2 or Web 1 fashion in which I am, as a user, just going with my web browser into https: domain, and some HTML is being loaded on my computer.

Specifically, Badger, like many other -- many other projects in the blockchain and, generally speaking, in tech, and I think probably in the IT industry, used a service called Cloudflare. Cloudflare is a very popular CDN and anti-DDoS service similar to Akamai for those who are familiar with. And this is essentially a service which is a centralized service that is sitting in front of the website and allows them to protect themselves from DDoS attacks and also to accelerate content to their users.

What happened to Badger, if we go to the next slide, which is quite interesting, is that at some point, their credentials to control their Cloudflare account were compromised. So basically there was a group of hackers that were able to go and modify the configuration of Cloudflare for specifically the Badger interface. And what was happening, if you were the user of Badger, right, you wanted to deploy your bitcoin over there, you -- instead of you kind of loading into your web browser the content, the HTML pages from Badger, the content was manipulated by the hackers that were sitting in between you and Badger on the Cloudflare layer, right?

So what the hackers cleverly were able to do is that they were able to insert code that manipulated you as a user to sign a transaction that pre-authorized what we call an approved transaction in DeFi. It basically pre-authorized your wallet to send assets in a future date to the -- to an address that is controlled by the hackers and the attackers. And they actually orchestrated this attack for about 2 weeks going unnoticed. So they created the pre-authorization across many, many different wallets, and then in a single day, they basically withdrew all the funds that were in that wallet -- in those wallets, and they were able to harvest $120 million worth of assets.

That's basically a pretty interesting attack because what it actually did, it basically diverted transactions from the users, and, as probably most people here already familiar with, there is no resource, right, for most of those funds, and, therefore, the assets were gone to the hackers' wallets. And that goes actually to the point that Dan made is, like, it's not sufficient always to deploy fraud controls that are operating, like, in traditional finance where you can analyze what is going on and then maybe, like, you know, try to reverse some transactions. You actually need to eliminate some of those vulnerabilities from the very beginning.

Now, if we skip to the -- to the next slide, I want to probably focus on maybe the most complicated part of DeFi security, and that's basically attacks that are operating on the smart contract level itself.

So another DeFi protocol, called Euler Finance, they have a pretty broad, I think, broad suite of services. They allow you to decentralize trading. They allow you to do decentralized lending and borrowing. And one of the interesting things is that they allow you to basically trade 10x -- basically a 10x leverage on the collateral that you deposit. So without going into the specifics of how their protocol actually operates --

If we go to the next slide --

-- what happened a week ago on March 13th or 14th, was that they had a logical bug in terms of how they were attributing collateral versus depth, right? And through a sequence of fairly complicated financial transactions that included the flash loans, leverage borrowing, and some deposits back into the protocol, the attacker was able to create a mismatch between what were the actual -- what was basically the leverage vis-a-vis the collateral that was in the protocol, and, therefore, were able to trigger a liquidation, but that liquidation not only basically took their position. That liquidation had actually gone into all the positions of all the other users on the protocol, and they were essentially able to withdraw then most of the collateral that was sitting in that protocol, making a profit of $200 million, right?

And that was something which is not -- the real complicated thing over here is that this is actually not an attack that we are familiar with from, I would say, traditional IT cybersecurity. This sits somewhere in the middle of a cybersecurity issue and real financial manipulation of a protocol. And because those protocols are effectively operating like, you know, a financial contract, there is no way for you to reverse back the result. So that basically covers, I think, sort of the technical details of three different case studies.

If we go to the next slide.

And what I think that is interesting is that the industry is currently situated in a fairly mature state in terms of both key management and transaction security where there are good examples of how this is being done. And maybe in those areas, we are already at the maturity where we can put some policies and best practices. Definitely on the last example of smart contracts, it's still an ongoing research. There is a lot of debate of what are the best approaches to tackle the challenges over there, and over there, we just need to continue with the research, put some best practices -- my opinion -- and continue sort of evolving how security is being done. Thank you so much.

MS. HOUSE: Thank you so much, Dan and Michael. So TAC members, we have now heard about the significance and challenges related to decentralized finance and, more specifically, digital assets and blockchain technology. To further consider these important issues, is there a motion from the body to recommend to the Commission that it establish a committee on digital assets and blockchain technology?

(Moved.)

MS. HOUSE: So moved. Is there a second?

(Seconded.)

MS. HOUSE: Lots of seconds. Thank you.

It has been moved and properly seconded that the TAC establish a subcommittee on digital assets and blockchain technology.

Is there any discussion or any comments on the importance of this subcommittee and any potential topics that folks would like to make? I'll ask for any of -- any interventions to be very brief, but any comments folks would like? Todd, you first.

MR. CONKLIN: Thanks so much, Carole. Treasury this year has been extremely focused on making sure we're expanding our sector risk management remit to include the private sector. And with that in mind, we did -- we did do outreach, and Trail of Bits and Fireblocks participated with Treasury in a gathering of minds in California a couple of weeks ago. And we went through a deep dive of all of these incidents and added even more to the list, and every single one of the instances that we laid out had no direct connection to any vulnerabilities with blockchain. They were all general with cybersecurity vulnerabilities and exploits. That is true of any firm anywhere in the economy, and I just want to make that point clear.

So then the question becomes how much of the vulnerabilities are really just a matter of culture within startups broadly, which apply, of course, to this sector being it's relatively new. And how do we then come together to impact and support that -- those startups with the full leverage of the U.S. Government and all the information that we provide broadly across the whole sector, making sure that reaches these startups as well? So happy to participate in this going forward.

MS. HOUSE: Thank you, Todd. Before we come to a vote, any other remarks about specific areas that are worthy of the -- of the committee's attention? Underneath the subcommittee, I know I've listed off some of the different areas that have been identified, looking at the why of DeFi, what problem it's actually solving, looking at different applications of it, vulnerabilities, issues, the policy issues, the legal frameworks as well as the technologies that need to be noted and developed. If there's any other comments, then happily turn to you. Thank you so much.

MR. PALMER: Yeah, thank you. Real quick, I think, echoing what we've already said, in addition to that, I think not a one-size-fits-all for DeFi is really important. We've kind of talked about it very broadly, but there could be specific products or financial assets that this works for very well but maybe others that it doesn't. So I think about best X in securities. How do you -- how do you best X in a global financial product that's decentralized? There's an intermediary that does that for you today, so very key things that I think we need to really think about.

Also, what happens with front running market data. This is more of a technology issue, but DeFi and everything being on the blockchain inherently slows things down, right? So there's this -- now this concept of additional information front running how slow is the market data being provided to the public, whereas if you look at how modern financial markets are, at least in the U.S., that's very not instantaneous but pretty close. And the markets that have been operating in these have been building that purposefully for that. So it's a big topic, I think, that should be covered as well in subcommittee. Thank you.

MS. HOUSE: Thank you. I appreciate that. We'll move to those on the phone. I see that Jennifer would like to make a comment.

MS. ILKIW: Can you hear me?

MS. HOUSE: Yes.

MS. ILKIW: Perfect. Oh, hold on. I'm staring at myself the way the -- I've set the view up. Hold on. There we go.

So I think when listening to everybody's comments, some of the words that I really picked out were "regulation," "compliance," "governance," "identification," "accountability," "controls," "policy." I think when we look at the traditional, I guess, centralized markets, I think what a lot of people forget is that these markets have developed over 200, 300 years. There's been a huge amount of innovation within these markets. This is not a market that has stood still. So I think as the committee looks, they really have to look at what we've also done in centralized finance to focus on risk management, the focus on financial stability, the focus on investor protection.

So when we look at DeFi, all those things have to be paramount as we're looking at helping to develop these markets, and to grow them, and to make sure that they fit how the markets, how people work, how financial markets work, and how it works within the traditional and the more innovative market space.

MS. HOUSE: Thank you, Jennifer. Michael, if you'll give closing remarks before we move to a vote.

MR. GREENWALD: Yes, very briefly. I just think for this subcommittee, there should be a focus on economic competitiveness and that theme throughout in addition to everything else you mentioned.

MS. HOUSE: Thank you, Michael. If there's no further discussion, we will now take a vote on the motion to establish a subcommittee on digital assets and blockchain technology. As a point of order, a simple majority vote of the present TAC members is necessary for the motion to pass.

For those in person, could I please see a show of hands for those voting aye.

(Hands raised.)

MS. HOUSE: Thank you. A show of hands for those voting nay.

(No response.)

MS. HOUSE: Thank you. Now if we can move to those participating virtually, please indicate "aye," "nay," or "abstain."

(A chorus of ayes.)

MS. HOUSE: Thank you. So noted. The ayes have it. We will submit the necessary paperwork to the Commission to establish the subcommittee, and we will be seeking TAC members to serve on the subcommittee.

Thank you all so much. Appreciate your patience with us going over a little bit. We'll now take a 10-minute break and reconvene at -- is that -- is that -- sorry --

SPEAKER: Two-forty.

MS. HOUSE: Two-forty. Two-forty. Thank you.

(Break.)

MS. HOUSE: Welcome back, everyone. We are ready to explore our second broad topic of the day, Ensuring Cyber Resilience in Financial Markets. To begin the discussion, our first presenter will be Todd Conklin, deputy assistant secretary, Office of Cybersecurity and Critical Infrastructure Protection, at the U.S. Department of Treasury. Todd will present on Treasury's Office of Cybersecurity and Critical Infrastructure Protections, or OCCIP's efforts to support sector resilience. Turn it over to you, Todd.

MR. CONKLIN: Okay. Thanks so much, Carole. It's great to be here, and it's great to be part of this Commission. I'm looking forward to the work to come. So first, I'll start with a very brief overview of Treasury's OCCIP, the Office of Cybersecurity and Critical Infrastructure Protection. I think I have some slides as well that we're going to launch through.

So generally, OCCIP is responsible for sector risk management of the financial sector, and I'm going to cover a couple of different initiatives that we launched over the last few months, one which is a Treasury cloud study and report which we publicly released a few weeks ago. And there's going to be a series of follow-up actions that we're going to work with the broader sector to close some of the gaps that that report identified. Treasury also worked with, Carole, you at one point, and the NSC team, and CISA to rebuild our cyber incident communications and Cyber Incident Response Playbook in the lead-up to the Russian invasion of Ukraine. And actually, the first time that we leveraged that playbook was during the ION incident from a few weeks ago, so there's some lessons learned from that as well that I will talk through.

So I think I have about 25 minutes. I'll leave some time for questions towards the end, so I'll try to blow through some of the meteor slides, but feel free to read through them as well.

So as sector risk management agency for the financial sector, Treasury's main goal is to ensure that the U.S. maintains the world's most secure and resilient financial system by spearheading a whole-of-government efforts to increase the cybersecurity and resilience of the American financial system. And we've got a lot of mature structures for the traditional financial sector. In particular, there's a Financial Services Coordinating Council, which is a formal structure that was established by the -- a presidential working group memorandum from several administrations ago. So we now have a 20-year playbook where the private sector formally engages with the Treasury Department and is able to share cyber information, and also general concerns, and also work with us on policy development.

So in addition to that group, Treasury also chairs the G7 Cyber Experts Group where we attempt to drive international norms and policies across the G7 countries. We co-chair that with the Bank of England. And also Treasury chairs the Financial and Banking Information Infrastructure Committee, or FBIIC, which is where all of the Federal financial banking regulators get together to discuss critical infrastructure and cybersecurity issues and drive policy normalization through that -- through that group. And the FSCC that I mentioned also has direct lanes to plug into the private sector. It has direct lanes that plug into that group through that formal apparatus, which has been, again, been in place for the last 20 years.

Additionally, Treasury, through its -- through its Intelligence Office, does rapid declassification of pertinent cybersecurity information for the sector. And if there ever is any specific intelligence pointing to any vulnerabilities targeting any one firm, Treasury does its best to, as close to real time as possible, get that information over to the firm either in a cleared way, or we try to declassify it in the cases where there aren't cleared personnel within a potentially targeted firm.

And then finally, Treasury administers the Hamilton Exercise Program. This year alone, we have over 12 exercises that we're going to work, along with our sector participants, to try to identify vulnerabilities within the financial sector, critical functions. So that's, again, another longstanding program that Treasury has implemented but probably isn't very publicly known, so we're trying to do a little better job of branding that, especially for some of the newer entries into the -- into the sector.

So if we go to the next slide.

So I alluded to the FBIIC already. So this is, again, where all the Federal financial banking regulators, including CFTC, meet Treasury to discuss cybersecurity and all hazards issues. So the senior leaders meet quarterly, and that's Deputy Secretary Adeyemo, Secretary Yellen, and then all the heads of each of the Federal financial banking agencies, to discuss FBIIC activities. The senior leaders at the start of last year requested that the FBIIC take on cloud adoption across the financial sector as a potential issue. So Treasury, with its FBIIC partners, began to develop a consultation network with more than 50 financial firms, academics, think tanks, cloud service providers, to really try to understand where the financial sector is currently in its state of cloud adoption.

So if we go to the next slide.

And we released a, which is now a fully-public report, a few weeks ago and did a major outreach effort to all of the -- all the firms we interviewed in advance, in addition to the many of the cloud providers. So the top line is that, and not a surprise to this group, I'm sure, but cloud service is -- really is no longer an emerging technology within the financial services sector. It's widely used for what we call software as a service, so video email, and video conferencing, and communications nearly across every single financial firm.

That being said, there is still a very fairly limited use across what we call infrastructure as a service. So critical assets, critical financial banking infrastructure amongst the major financial firms, it's fairly limited at this point. That being said, what our interviews revealed is that many of the larger financial institutions have a three- to five-year adoption strategy for which there they're going to layer in some of their more critical assets on some element of cloud.

The story is much different when we expanded the interview list beyond the critical infrastructure, larger financial institutions, and global financial institutions to the local and community banks. The story was much different in that the local and community banks felt so much pressure from fintechs, and, additionally, their third-party vendors moved to the cloud without the decision-making process of the actual C-suite of the local and community banks, right? So local community banks are stuck in a position where they have to go to cloud whether or not they want to.

So with the larger financial institutions that have this ability to have this three- to five-year road map, a lot of our local and community banks are now 100 percent cloud, and they don't necessarily have the talent at their disposal to implement the shared security model that cloud requires, right? So it's a more acute problem for our local and community bank partners who are -- who have historically been reliant on third-party technology providers for a lot of their functions and services. So there's really a tale of two stories to compare -- when you compare the larger banks to the smaller institutions, and we all know the challenges with talent acquisition generally in the cybersecurity space, obviously much more acute on the smaller financial institution side than larger side.

If we could go to the next slide.

One note at the top. So we did add to the report Treasury's own the cloud adoption strategy. As Carole is aware, I was one of the CIOs at Treasury that was really focused on cloud adoption. So we, through our national security apparatus, we started adopting cloud about five years ago, and we really have the cloud first mindset for a lot of -- a lot of our workflows now, which is much different from where we were are maybe five to six years ago. So Treasury is also a user of cloud obviously, and we did layer our anecdotes and notes from our own cloud adoption strategy into an annex of this report, if you are interested in learning more about Treasury's own cloud adoption, pitfalls, and success stories.

So anyway, generally, the potential benefits of cloud, why is this even a discussion? Why are firms interested in even moving some of their core infrastructure and critical assets into the cloud? Generally, everyone seems to agree on the big three: redundancy, scalability, security. So cloud -- when implemented properly, cloud services offer physical redundancy with the potential to operate from multiple availability zones, which are physically or logically isolated data centers that hold -- host cloud services, right? So if you're talking about a local community bank, you just can't compete with that level of a availability, right, that the cloud service potentially can provide when implemented properly.

That being said, the multiple regional availability model is also much more expensive to operate as opposed to a single-region approach within cloud. So the other piece of that then is scalability, and that gets to the competition element, with fintechs in particular, that a lot of firms see the access to scalability that cloud offers them as an opportunity to be much more competitive in the marketplace. And again, that's much more acute for the local and community banks scenarios and in security. You have several large firms that spend billions of dollars in cloud infrastructure support that, again, the local and community banks, it's hard for them to compete with that level of investment and that level of technical aptitude that the larger cloud providers put into their security offerings.

That being said, if we go to the next slide, and the security argument.

I've seen some -- I saw some heads nod, and some say "no" when I mentioned that piece. Obviously, there's another side to that concentration piece as well that could also be a negative, which we get into in our six main challenges but that the report identifies. So while the report went into great detail on talking about the benefits of cloud adoption, we also then unpacked six core issues that the firms' message to -- some of them was around -- and this is really the number one -- around transparency. And this is why some firms are not actually going to invest much more in cloud because they do still have some concerns that they lack the information necessary to conduct due diligence and monitoring of the cloud providers, and that's not universal. That came out from just a limited number of interviews.

Additionally, there's gaps in expertise, in tools. I've already alluded to the -- this being much more of an issue for the -- for the local and community banks based -- still a challenge, though, even for our largest financial institutions, who are having difficulty staying competitive with the Netflixes of the world, right, for, one, cloud talent, two, cybersecurity talent as well. And then third being exposure to potential operational incidents, including from incidents originating at cloud service providers themselves, and that's where obviously that concentration becomes a negative on the security side potentially. Fourth challenge being potential impact of market concentration on the sector's resilience. Fifth being dynamics in contracts negotiations, and nearly every firm we spoke with talked about pain points in their first contract cycle with cloud vendors, and some of the larger institutions highlighted stories of if we knew -- if we knew then what we know now, right, how much different our first sets of contracts would have been.

Again, this is an issue that -- where it becomes much more challenging for the smaller financial institution who don't have those robust legal teams to negotiate with the -- with the larger cloud providers in particular. So there is a bit of a disparity there in some of the contracts that we observed for some of the smaller institutions, even things like not negotiating upfront access to all of their keys and all of their data in the event that they did want to extract all their information from one particular cloud provider, right? Things that have become standard across some of our larger and institutions haven't quite become standard yet across our smaller.

So, and then, of course, there's the international landscape and broad international regulatory fragmentation issues. DORA, of course, is top of everyone's mind these days, so there's just a lot of international regulation coming out, some of which is not completely coordinated through things like the G7, which I alluded to earlier.

So those are the positives. Those are the top six negatives. What are we going to do about it?

So if we go to the next slide.

So Treasury has established a strategic vision for supporting the resilience of the financial sectors' use of cloud services. Treasury is positioning itself now to take a leadership role in this space and begin the process of making sure that we're making cloud as secure as it possibly can be for our entire financial sector, including all participants.

So if we go to the next slide.

We're going to establish a Cloud Services Steering Group. It's going to be led by leaders from the FBIIC, FSOC, and also Treasury. And additionally, we're going to have a partner group that's going to be at the CEO level of financial institutions, which are going to plug into this executive steering group, and we're going to work on a series of very specific items throughout the remainder of this year.

One is development of common definitions and terms to make sure, one, that we're all using the same lexicon and terminology. And this came out even as we were drafting the report when we were trying to synthesize all the edits and inputs that we got from the Federal financial banking regulators plus Treasury's own CIO teams, things like multi hybrid cloud, hybrid cloud. Everyone was using different definitions for kind of each nuanced area, so just coming up with one common lexicon that then each of the financial banking regulators can then take to the onsite exams, and then we could potentially even share exams across the entire Federal regulatory landscape. It's potentially going to close -- just that one simple lexicon workflow might potentially close quite a few of the -- of the gaps that that we observed in the report.

Additionally, we're going to explore the authorities required to provide more direct oversight of cloud service provider infrastructure itself, so not just cloud service and infrastructure through the lens of financial services firms but actually by examination potentially of the cloud service provider. So that's a workflow that we're going to begin in earnest in the -- in the coming weeks.

And then on the private sector side, one thing that we ask them to help us lead is the contracts piece being that they are the core customers of the cloud service providers. We're going to try to leverage it --

Go to the next slide, please.

We're going to try to leverage a lot of the work that SIFMA already began in terms of contracts best practices and making sure we could scale that across the whole sector beyond just the securities grouping that SIFMA is most focused on. Additionally, we're going to ask that the Cyber Risk Institute work with our NIST partners to establish some very clear NIST-centric frameworks that the entire financial community could then leverage for their cloud adoption strategies. That way we're helping, as best we can, the local community banks that don't quite have the talent access, that we're giving them clear roadmaps for their cloud adoption strategy, leveraging the know-how and expertise of the sector participants who have already -- who are part of the Cyber Risk Institute and who have already adopted cloud themselves, some quite successfully.

So if we go to the next slide.

So that's -- that concludes the cloud piece, so I can stop there and take some questions, and then I'll -- I can kind of switch into incident response.

MS. HOUSE: Go ahead, Hilary, and we've got time for questions also after Kevin's presentation, but, Hilary, let's go ahead and kick this off to you.

MS. ALLEN: So this is just a question. It's related to concentration risk. But if you have multiple banks using the same cloud at the same time, is there or has any thought been given to how you might have to stagger them after a problem has occurred? So if everybody's trying to download at the same time, that could potentially tank the cloud. So I was just wondering if any policy consideration has been given to the systemic consequences of everybody trying to download at the same time and whether there's sort of any plan for staggering or anything like that.

MR. CONKLIN: Yeah, a great question, and maybe I'm -- I might punt that question to the next piece, which is the incident response, which is inserting Treasury at the -- really at the center of public communications during an incident, which is really kind of a new model than what we previously run. So I could -- I could envision that touching on your -- on your question a bit more deeply.

The other angle of that is that we're really --

we're really just thrilled that we have the full commitment of the cloud providers themselves to participate in these -- in these workflows. And I think one of the -- one of the concerns going into it is that cloud providers would encourage lock in, right, for their own, right, cloud offering. And in the -- in the coalitions that we've been able to develop through the creation of this report, I'm optimistic there is a potential future where we do have a more realistic kind of hybrid multi-cloud where all -- that being said a lot of what the technical experts implementing cloud right now will say is that it's not really viable right now to run this hybrid cloud model. It's not technically feasible at this point at the scale we would need it.

That being said, I think there's a lot of work being done right now, that we're going to be in different -- a different place five years from now. And I think that will -- that will help resolve some of that -- in addition to the multi-regional option, that will -- I think the multi-cloud hybrid approach will ultimately also help as a -- as a backstop to --

with concentration. And generally, we're going to take on concentration risk as a topic by itself next year once we kind of get through these basic items first. So thank you. Great question.

MS. HOUSE: Thank you. Stanley, you have a question for Todd.

MR. GUZIK: So, Todd, thank you. So I've been implementing and building software applications in the cloud probably now for over 12 years, migrating legacy systems to the cloud. So the question I have is on, you know, the work with the small and the regional banks. How much of that -- you know, the negatives or the challenges are very legacy applications or a technology. So they're running on mainframes, which is really challenging, you know, porting a mainframe to the cloud. They're on -- running in AS/400, or they have client server, like desktop software, connecting to client servers because as you move to the cloud, there's a significant amount of refactoring. You need to -- you know, it needs applications that enable them to move to the cloud.

So how much of that is, like, the negative, you know, as part of the -- is it a top seven negative?

MR. CONKLIN: I would probably even say it's higher. Even at Treasury, that's -- for a lot of Treasury's largest bureaus, some -- one of which does have a mainframe application rate, I think it's just very, very ineffective from a cost perspective for Treasury, for example, right? So I can't imagine having to pitch that then to a board and then layer on the aspect of not having talent to navigate through that, so then you're reliant on a third-party contract service coming in then to help you implement that transition as the local and community banks. So it's just a significant -- it's a significant problem.

But many -- that being said, many of the local banks already are outsourcing their infrastructure anyway, so it's not necessarily a case of their own legacy infrastructure. It's just that they've already outsourced their IT, say, to a Kaseya-type firm, and Kaseya makes the decision that they're no longer going to have their legacy, you know, on-prem offering. They're going to be 100 percent cloud, and then -- and some of those then were offerings that were private on-prem that they set up for some of these local banks, which then they basically --

I'm not picking on Kaseya. Kaseya didn't do this. I'm using Kaseya as a hypothetical, but then you can imagine a scenario where that third-party vendor says we're not going to maintain this on-prem offering for you anymore, you have to go to our cloud offering, and then the local bank doesn't have cloud experts at its disposal, so. Okay.

MR. SIRER: I just had a quick question about disclosure requirements. I know that encrypt, everything is out in public. Everything is transparent. We get to find out about all the hack. In the banking universe, when a bank does not lose PII but does lose financial, you know, assets, but it does lose money, is it obligated to reveal to the public what happened? Do we get to find out as regular citizens when banks lose significant sums of money to hacks?

MR. CONKLIN: So that that was actually work stream that the FBIIC took on last year that we completed work, and both angles of that work being that if there -- if there was consumer data on the U.S. Government where there was a breach, the U.S. Government committed to a very, very expedited notification to the owner of that data. So that was part of that FBIIC process.

That being said, I'll defer your questions to the SEC guidelines, which I know there's some information out there now. SEC is taking a leadership role in this space, in particular around data breach notification and timelines. So there's -- there are some updates to that, which are currently up for discussion publicly. So if you did have an interest in exploring that issue, you could -- you can go to the SEC website and submit a response to their request for information on that. But there should be updated rule on that, I would guess, fairly soon, so.

MS. HOUSE: The Federal banking agencies also have incident reporting requirements underneath some rules that they published recently, but I don't -- I don't know the extent to which that included notifying the victims versus notifying the regulators. So, again, as Todd mentioned, defer to the regulators but just wanted to point that there's also rules, not just the SEC's, for public companies obviously, and those under their jurisdiction but also the banking agencies. Thanks, Todd. Do you want to take us through the rest of your presentation?

MR. CONKLIN: Okay. Let's go. All right. So now, we're going to we're going to switch from cloud to general incident response. And Treasury worked with the White House, and CISA, and broadly DHS at the start of the Russian invasion of the Ukraine to make sure that our instant response playbook was calibrated properly for the -- for the potential notional thought of a nation-state actor potentially trying to impact some element of U.S. critical infrastructure, and the question being how much does that cause us to adjust our incident response playbooks to potentially contemplate for a higher severity level of incident than we -- than we maybe normally would have potentially oriented our incident response towards.

So the White House, and much -- with much credit to Cal's leadership, the White House stood up a Unified Coordination Group on February 22nd of 2022 to make sure that the U.S. Government as a whole was starting to think through a lot of these issues. And through that UCG --

If we go to the next slide --

-- we established a playbook that contemplates really three levels of escalatory activity that a nation-state may potentially try to take action, that a nation-state may try to impact in the event of escalations of hostilities. And we broadened it beyond just one potential notional nation-state to include any talented nation-state adversary that might want to impact harm on our critical infrastructure.

The actual nuances of those levels are confidential, so I'm just -- I just have a numbers scheme on this chart. But what we did for the financial sector where we already had a five-level incident on schema, so five levels of severity potentially impacting the financial sector as a result of a -- of a cyber incident, we overlaid the nation-state layer onto the FBIIC incident scheme. So that's what's displayed here.

And then on the next slide, what we did, once we had that kind of core baseline established, we went out to the financial sector and said, okay, for each level of incident, what do you need from Treasury. If you were impacted at this level, what do you need the Treasury Department to do? What do you need the Federal banking regulator agencies to do to help you navigate out of this particular incident or issue? And the -- and the full list of inputs that we got from the financial sector are on display in this chart. But the number one thing we heard, and this goes all the way up to the CEO level, was we need the Treasury Department to maintain clear and direct channels of communications between the United States Government and the firms, and also control the public relations response to an incident.

And this is critical because after we started the process of constructing this, and Ari alluded to the Colonial Pipeline incident, Colonial Pipeline happened, which was a ransomware attack impacting -- it was really a -- what I would consider a lower-level cyberattack impacting a firm. And we went into the weekend at the start of that impact cycle not thinking there was going to be a supply chain issue. And within two days, we had lines around the block up and down the East Coast because people were concerned they weren't going to be able to get their gas, and that human behavioral response is what caused the supply chain issue ultimately, right? So there is this -- a low-level cyber issue that really should've been isolated. Human behavior takes over, and that causes the supply chain response.

That's exactly what we're trying to avoid in the financial sector in the event that there is a lower-level cyber impact targeting one of our critical firms. We don't want a situation where it spirals into some sort of supply-chain-type crisis or bank-run-type crisis that we observed during the Colonial Pipeline situation. That's what this updated playbook is oriented towards. So we actually -- we completed it in June of last year. We ran an exercise with the sector in September, and we got to deploy this playbook for the first time during the ION incident.

So if we could go to the next slide.

For those of you not familiar -- I'm sure most of you are -- I'll just do a really, really high-level overview of ION. So ION was impacted by a lock-bit ransomware attack sometime around overnight on January 30th into the 31st. The U.S. Treasury became aware of the issue around the afternoon of the 31st with really not much clarity on exactly what was impacted. There wasn't much information at this point coming out of ION itself. Treasury started to get outreach from Ireland, in particular, and then soon after Japan and Bank of England indicating that there were some significant delays in derivatives processing. So this was all into the evening of the 31st.

And our preliminary assessments with our international partners was that ION, which is a significant market player offering third-party vendor software in the derivatives and also trade space, they have a significant footprint around our central banks especially. So global central banks leverage them for quite a few different software applications. They also have broad market leadership positioning in the in the Treasury space, less so in the derivatives space.

And over the course of the last few months, they've been on a bit of an acquisition spree. So you have this potential sprawling impact zone for a -- for a firm that, what we found later, many institutions didn't even unclassified necessarily as a -- as a critical third party vendor, right? So many firms who onboarded ION didn't use the highest level of scrutiny that they use for their most critical third-party vendors. So I'm painting a picture that we had a very heightened concern going to sleep on January 31st.

And then if we go to the next slide.

So we woke up on February 1st really a complete unknown in terms of the number and type of ION services disrupted, unknown in the number and size of financial institutions that were impacted, and unknown for the size of outstanding debt held by impacted traders and size of creditors. All we knew -- and by this point, by the time we woke up, Japan had completely disconnected from ION, so the situation seemed to be spiraling in the -- in the wrong direction that morning.

Very, very, very quickly, the SEC and CFTC took a leadership role here, and SEC, in conjunction with CFTC, was able to ascertain the exact software impacts at ION, and, fortunately, it wound up being limited to about 11 of their applications, most of which was in the derivatives market. So the smallest market share element of ION was the one that was ultimately impacted. CFTC then, through a number of engagements -- I don't mean to speak for the commissioner here -- through a number of engagements with the -- with the sector and the firm itself now on February 1st, confirmed that the impact was limited to roughly 41, 42 financial firms, and that there would -- there was no significant impact to our central banks, in particular.

So within a matter of hours, we were -- we were able to basically get a really clear operating picture. So that concern that started in the morning, by the afternoon, it was clear that we had a much -- a much less severe situation, and throughout that process, we convened the FBIIC. So we brought the FBIIC together. We brought the FSCC into the -- into the fold as well and got a clear common operating picture across the entire Federal banking regulatory agencies plus the private sector, which included SIFMA. It included the FS-ISAC. It included the Analysis and Resilience Center for Systemic Risks. So all the critical financial sector firms all were on calls throughout this multi-hour period of heightened concern to make sure that we were all on the same page. And we were all clear by the afternoon that this was not a systemic issue.

So that being said, if we go to the next slide.

So as we clarified our view of the situation being less severe than we thought from the day before, the media started to get wind of the issue, and they were taking the approach from where we were the day before of ION is a significant player. This is going to be a systemic issue. There's already -- some regions have already disconnected completely, that this is going to be broad chaos. So we started to see some articles come out in the press around noon on February 1st, right around the same time we made the broad FBIIC assessment that this was a lower-level one incident.

Treasury was then able to activate the public communications playbook and did direct outreach to all of the reporters who published articles on it, and asked them to update their stories with Treasury's assessment that the FBIIC had convened and that this was not a systemic issue. And we also issued some proactive statements, which were then carried by some other news sources. So within a matter of hours, we were able to adjust the media narrative from one that was extreme concern to one that -- not a good situation. Some firms were impacted, but the situation is not systemic and it's under control.

So a lot of lessons learned being that that was the first time we did leverage that process, but it did -- it did work to avoid the Colonial Pipeline type situation that we designed this process for. So I will stop there, and --

MS. HOUSE: Thank you so much, Todd. Really appreciate that overview. So, Dan, I do see that you have a question. I'll make sure that you can open up our discussion right after Kevin's presentation. So for our second presentation regarding cybersecurity issues, we have a presentation from Kevin Stine, chief of the Applied Security Division at NIST Information Technology Laboratory at the National Institute of Standards and Technology, one of my favorite agencies. Kevin will present regarding managing cybersecurity risks. Kevin, over to you.

MR. STINE: Perfect. Thanks, Carole. I appreciate it, and thank you to the Commission for including us today. I definitely appreciate Todd's comments. We were fortunate to have had and continue to have a very strong relationship with our Treasury colleagues and, I would say, actually, the broader financial sector. In fact, you know, just last year we celebrated what we count as our 50th year in cybersecurity dating back to 1972 at NIST.

And prior to NIST, we were called the National Bureau of Standards, and our work in cybersecurity at that time really started with the development of the data encryption standard. Prior to that point, encryption was really a military-grade application, and the financial sector at that time identified the need for encryption to satisfy some of the business needs going back to the early 1970s. So we worked with the broader sector. We worked with companies, like IBM and a handful of others, certainly at that time to then develop and issue a standard for data encryption. A lot's changed over those 50 years, but I think we still do have a very strong relationship with the financial sector and many others as well, so just a little bit of history for you to take away today.

So, again, Kevin Stine, National Institute of Standards and Technology. If you're not familiar with NIST, we are a part of the U.S. Department of Commerce. We're a non-regulatory agency. Our mission, very simply put, is we seek to promote innovation and industrial competitiveness through standards and measurement science. I think NIST came up maybe in some different context over the course of the day today, so if you have any questions on other cybersecurity-related work, happy to answer those as we go forward.

From a cybersecurity and really, increasingly, a privacy perspective, we think about not just our mission but really our purpose, which is to cultivate trust in technology. We try to do that through better standards, better technology, better measurement science. And this idea of trust, you know, having that foundation of trust is really critical, really based on standards is critical to provide a consistent level playing field but also to provide kind of this platform for innovation. And there's a lot of innovation that I think happens within this community for sure, so we're excited to be on this journey with you.

Perfect. Next slide. Back up one. Sorry. There we go. Perfect. Thank you.

So one of the, I guess, a core tenet or a thread that we pull throughout everything that we do within our cybersecurity work in NIST is this notion of risk management. Look, every organization manages many different types of risks every day -- financial, reputational, operational, compliance, privacy, safety, you know, cybersecurity or information security. These are all managed each and every day, and I think frequently, we see these risks are managed in silos. I think there's a lot of challenges to what we think of as kind of the broader enterprise risk management, kind of this focus or function to kind of pull a lot of different -- diverse types of risks together under one umbrella, if you will, view those in the context of an overarching enterprise objective or set of objectives. Critically important there.

I think one of the -- some of the key points around enterprise risk management, certainly that risks can be managed, you know, in a means that kind of tie into mission impacts. ERM really helps to support more credible decision making on risk and opportunity information. Again, I say "opportunity" because, you know, risk can be both positive and negative, and there's opportunities that can be had there as well. I think one of the opportunities and really -- and a challenge as well is kind of this normalization of risks across the enterprise.

And one of the pieces that makes that particularly challenging is that, as I mentioned, many of these are managed in silos. Many of these have their own kind of language taxonomy. I know we're guilty of that in the cybersecurity space. We speak our own language that might not resonate with kind of broader enterprise risk folks or even folks in other domains. So that's certainly a challenge that we see, you know, beyond just the acronyms, just the specialty language that happens there.

So that's -- I want to focus in on that communications piece with, you know, my slide, and we can go to that next slide as well because I think that is one of the big challenges that makes cybersecurity more difficult to manage within its silo but also in the context of a broader enterprise or broader set of mission objectives.

And want to take us back real briefly to 2014 when we issued, again, based on an executive order that drove us in this direction, issued what we call now the NIST Cybersecurity Framework. And very simply put, think of it as a tool to help organizations better understand, communicate, manage, and reduce cybersecurity risks, standards-based tool to help do that. What it does is provide very much a common language, a more common and accessible language, to help organizations within the organization talk about cybersecurity risks. That could be from kind of, you know, the maybe overused phrase of the C-suite and the board of directors, to the bits and bytes folks in the data center, and everybody in between. It can also mean kind of your horizontal, you know, between your organization and your partners and suppliers being able to talk about cybersecurity risk, talk about requirements events and expectations, talk about your own capabilities from a cybersecurity perspective, but also being able to talk about cybersecurity and the things that your organizations do or provide with your customers or consumers of your products and services.

And maybe the fourth I would add is kind of an audience. And we see a lot of potential here is, particularly for those in heavily-regulated sectors or multi-regulatory environments, the ability to talk about cybersecurity with regulators or with kind of organizations that have some sort of oversight responsibilities for you, for example. I think there's a big benefit to that common language, especially in what we're hearing a lot about now in this area of regulatory alignment where many organizations fall within multi-regulatory environments. There is a finite list of things, if you will, from a cybersecurity perspective or a technology perspective today that can satisfy many of those types of requirements. So how can we best align the requirements, the language that we use to talk about those requirements but also the language we use to talk about how we demonstrate or articulate how we've chosen to take on those particular requirements and implement capabilities to address them?

A few other kind of points about the framework. Again, it's risk based, and it's -- and it's outcome based. Again, it's not prescriptive. It doesn't prescribe a specific control, or capability, or technology be used but rather focuses on the outcomes. Those outcomes could be very different from one organization to the next based on your risk tolerance, your prioritization of those external requirements. You might have expectations from partners and suppliers as well, but that outcome-based focus can allow you to have greater flexibility for how to achieve that outcome, based on your resourcing, your capability, kind of the availability of different tools and technologies that could be used to help you satisfy that capability.

It's meant to be paired. The framework is meant to be paired.

Let me back up from the mic a little bit.

You know, the framework is just that: it's a framework. Again, it's a collection of outcomes that were developed in close coordination and work with public and private sector stakeholders. It's going to give you a little bit more about the why and the what but not a whole lot about the how. So it's meant to be paired with other guidance, other resources that can be helpful to help an organization along their journey to achieve a particular outcome in some particular way.

It's meant to be an adaptable resource to many different types of technologies, life cycles, sectors and uses, again, for a more agnostic framework. We've been thrilled to see the uptake of it over the last 10 years in every critical infrastructure and well beyond that into, you know, all different sectors and segments of the economy. And it's -- perhaps most excitedly for me, you know, the international uptake has been tremendous. We're up to 10 foreign language translations of the framework now, and we've been excited to see the framework be adopted by many nations around the world really to serve as that basis, in some cases, for their national cybersecurity strategies and approaches as well. I think we're seeing its uptake, or at least leveraging it in some ways, in foreign regulatory environments as well, which, again, back to that communications, that common language capability is critically important.

So we can go to the next slide.

So the framework has been out for about 10 years. I guess we're actually at nine years now. We started the process in 2013, so 10 years ago we started. And we've been excited, again, with the uptake of the framework, but we recognize that it has to be updated to be -- to maintain currency with the ever-evolving technology landscape, the threat landscape, and be informed by, you know, different happenings, including uses of the framework by sectors, by organizations, and even by nations.

We last updated this in 2017 -- 2018, sorry -- 2018, and certainly a lot's changed since then. There are a few key areas that, you know, we've seen evolution, not only within the technology landscape, but really as organizations have continued to use the framework, organizations of all shapes and sizes across all different sectors. We've learned a lot. We've learned about some opportunities for improvement and areas for emphasis -- for greater emphasis. And I wanted to highlight three of those because I think they tie in with some of the -- kind of the broader agenda items and interests of this group as well.

So the first is this idea of, you know, cybersecurity governance. And I don't -- I don't want to leave you with the impression that governance was not included or considered in the current versions. It certainly is, but I think what we've seen is with greater board-level or executive-level interest in cybersecurity as an enterprise risk, and certainly increased attention, you know, at a national and even global stage on cybersecurity over the last several years increasingly, the idea of governance, cybersecurity governance, has kind of, you know, risen in prominence and importance. And it's, in part, how do we talk about cybersecurity in a way that will resonate with the broader organizational governance activities.

There's certainly that common language that we talked about, but what are the things that we need to provide as cybersecurity professionals, the pieces of information that we can provide to broader organizational governance functions that will use their language, be using tools that they're familiar with to be able to view cybersecurity risk alongside other dimensions of risk that they are chartered to have oversight over foreign organization. And certainly, that model can extend into sectors and even nations as well, so we expect to have a greater emphasis on governance within the update of what we're calling CSF 2.0.

The second piece is greater emphasis on the importance of cybersecurity supply chain risk management, again, an area where we, I will say, kind of dipped our toes in a little bit back in 2018 timeframe with the last update. But there's certainly been a tremendous amount of interest and work over the last several years, to the point where we want to incorporate more supply chain considerations and really key practices in cybersecurity supply chain risk management into the framework.

I would say our overarching objective for our work in cybersecurity and supply chain risk management is really one of visibility. How do we provide better tools and information, whether it's standards or guidelines, or tools and approaches, that kind of the broader community is developing to help an organization gain greater visibility into their supply chains, and the partners and suppliers you're working with, their security capabilities, maybe some of the gaps in their capabilities? How do you express your requirements and expectations, and how can they provide back to you, if you will, demonstrate that they can meet or achieve the requirements and expectations you've set?

Depending on where you are in the supply chain, you may be one of the partners and suppliers or the third parties, if you will, so I think we can all be somewhere on that spectrum and somewhere in that alignment there. You know, we think the language of the cybersecurity framework can be very helpful for organizations to both express those requirements but also to be able to assess those requirements as well on how an organization is achieving those.

The third area that we're going to increase our treatment of, and all of these I describe as potential changes because we're in -- we're going to be heading into a more robust public comment period. And we want to get a lot smarter from the community on how to address these and certainly what's the right level, you know, the Goldilocks approach of have we gone too far, have we not done enough, and how do we land in the right place. And that's certainly the case in the cybersecurity measurement and assessment space.

You know, for a measurement science organization like NIST, I mean, cybersecurity is a hard measurement problem. I think if it was easy, you know, we or someone else would have done that a long time ago. There are things we can certainly measure today, very technical things, bits and bytes of, you know -- you know, entropy for example, and cryptographic algorithms, and those types of things. But the composition problem, in my mind, is where a big challenge is. How do you take a lot of the bits and bytes things that you can measure and begin to roll those up into some of the more qualitative measures, like am I more secure today than I was yesterday, or if I give you $10 in cybersecurity spend today, is there going to be a greater return on investment in terms of cybersecurity capabilities or practices tomorrow.

Those are the types of questions we understand organizations are asking, and not that we have answers, but we want to understand how tools, like the framework, can be improved to help provide greater information or greater approaches to help organizations come closer to being able to answer those types of questions. So the measurement assessment is very much a practical here and now, but it's also very much a research opportunity as well, so a core area for us to focus on at NIST.

Over the next month or so -- a month or a couple of months, I should say -- we'll be putting out some more draft materials related to the framework to really solicit more public comment in these and other areas. So we certainly in -- are our excited to get anyone's feedback from your organizations on how we can better improve the framework to bake -- make it a more useful and actionable tool for helping organizations to better manage cybersecurity risk.

Go to the last slide, I believe.

You know, I've flagged a couple of resources in passing as I was talking, if you have access to these slides electronically. I know we've all been trained to not click on the links. You're welcome to click on the links in the slide. Trust me. You know, I come from NIST. You know, cultivate trust. Click the link, but those are direct links to some of the resources that I talked about both in the cybersecurity and the -- in the cybersecurity supply chain space. Definitely welcome your feedback and involvement on any of these resources and really anything that we produce from the cybersecurity and privacy perspective and certainly welcome any questions that you might have now. So thank you/

MS. HOUSE: Thank you so much, Kevin, and I'll take a moment to foot stomp why, amongst many other reasons, NIST matters to you. Some of the discussion earlier about concentration of vendors, why the SCRM -- supply chain risk management -- work that's under way, and the great guidance, and everything else coming out from NIST, as well as other interagency partners is so helpful because only through illuminating your supply chain, identifying those points of concentration can you potentially understand the possibly devastating consequences that you can have when there is an aggregation or concentration of certain services across, whether it's traditional financial institutions or fintech.

And then also another example that I loved, since Todd brought up the ransomware example, which is so relevant here since it is both a cybercrime and a financial crime that is laundered through the financial sector and often through DeFi. I really loved -- NIST, you guys published a ransomware threat profile. I remember when it came out for comments. I can't recall if it was finalized, but they went through and identified the most common and prevalent vectors for compromise, in, like, the left-hand column and then the right-hand column identified the controls that NIST had previously published and how they mapped against being able to defend against those kinds of -- those kinds of exploitations and threat vectors. So those kinds of tools are just so critical for financial institutions to be -- to use to defend against the kinds of the kinds of incidents that I know Todd spoke to.

So at this time, I would like to open the floor to questions and comments from the TAC members. If anyone has any questions for Todd or for Kevin, please raise your flag. Oh, yeah, of course, Ben. Sorry. I forgot that you had raised it earlier. Apologies. Over to Ben.

MR. MILNE: Thanks so much. This question is more for Todd on the work that you're doing. As you were talking, the words that really stuck out to me were "incident response," "early warning." And, you know, one of the challenges with non-cloud systems is typically reporting, particularly SAR reporting, as it relates to incidents with information security, is typically not really being reported at the rate the crime is being committed. However, if the constitutions are based on cloud systems, reporting could presumably get much faster. And I was just curious how there might be some overlap in improved or even programmatic SAR reporting with some of the work that you're doing.

MR. CONKLIN: That's a -- that's a great question. Frankly, the SAR aspect of it hasn't been part of the OCCIP work, but it's a great flag. It's maybe something we can make part of this committee or something that will -- I'll definitely take back. But that's a -- that's a really great flag, and I think there's an opportunity there, too, as we on board the conversation with the cloud providers to see if there's any opportunities there. So good flag.

MS. HOUSE: Stanley, do you have a question?

MR. GUZIK: Yeah, just a comment to Kevin. So I'm a big fan of the NIST framework, so thank you for all that work there. Fully support and believe on the third-party risk, you know, the large amount of work that we normally do is reacting to vulnerabilities and third-party middleware software, whether it's, like, Log4j, Solarwinds, that's constantly happening.

But the comment on the enterprise risk management, especially reporting out to ERMCs, you know, reporting out to the board or the sub portions of the board, when cyber risk is reported, it's always consistently the same. It's high or it's elevated. It's elevated. But how do you actually -- like, in the frameworks about -- how do you actually measure -- you mentioned about the measurement. A lot of these frameworks or, you know, reporting out, even if you -- you know, the technology team is improving, constantly improving, it's always elevated. It's always high. It's, like, getting a little bit more of that granularity, yes, it will always be high, but measuring and showing the improvements. And do we get to a point where it's not high?

MR. STINE: I can't say that today. Yes, this is a big challenge area. I think certainly the approach we're trying to take is, are there better ways or more effective ways to measure cybersecurity capability so that we can have those types of measurements to have more granularity, and, you know, when we say "high," this is really what it means. I think a lot -- that, in my mind, would help really try to get to not just, yes, it's high, but what is the impact of it being high to the organization. How can we either quantitatively or qualitatively provide a little bit more context around the ultimate impact of those different -- of the risks that have -- that kind of bubble up to that enterprise risk level.

MS. HOUSE: Commissioner Goldsmith Romero?

COMMISSIONER GOLDSMITH ROMERO: Todd, I had a question for you on ION markets. One of the things you were saying that several financial institutions had not even listed them as a critical third-party service provider. So what is the lesson to be learned from that in how -- in how financial institutions or others would categorize their supply chain or the third-party vendors?

MR. CONKLIN: Great question, and I talked to a couple of chief risk officers from some of the G-SIBs, specifically, about that that gap. And it's an area where clearly the government can help provide some additional to the NIST framework kind of approach, provides some additional framework, more around risk management, which it hasn't been an area that really OCCIP, in particular, at Treasury has focused on. We've been in the incident runs and information-sharing piece of it, but how do we -- how do we help the sector with risk modeling broadly going forward?

And one of the projects we kicked off this year with the FBIIC is, we call it the SECURE Project. But it goes around all of the different third-party entanglements that the sector has. How do we begin to kind of shine a light on the more critical nodes of that so that the largest firms that have thousands of vendors can triage the third-party risk management onboarding process a little bit better? So how do we -- how do we add that intel mindset to the risk management space in ways that we haven't before? And it's really -- I think we're trying to kind of go down a new lane with that this year with the help of CFTC and the broader feedback team, so.

MS. HOUSE: Justin.

MR. SLAUGHTER: Thanks for that. Yeah, to respond to Stanley's point, I remember being told six or seven years ago when I was at CFTC that, you know, if you look at a lot of the cybersecurity world, it's like a soccer match where the score is 270 to 271. Basically, the problem is we are much worse at creating defenses than attacks across the board. So that -- it's not that it's been high because it's always high. The tech to defend is constantly a step back behind the kind of attack.

The number one thing we've seen, I think, both a paradigm and in general in my career, is this is where you need white hat hackers more than anything else, and you basically have to be constantly battle testing. The fallacy is thinking this is static, it's a dynamic risk, and you basically have to always go after it.

MR. STINE: Yeah, I think that's right, and I think that's why we -- every organization be -- should be thinking about this, not just in terms of in the -- in the language or using the framework, you know, respond and recover, but really resilient. And that starts at an enterprise level, you know, understanding what your risk tolerance is, and then being able to architect, have a resilient architecture that can withstand or continue to operate in light of kind of, you know, the challenges that you're facing. So I'd certainly agree to that.

MS. HOUSE: Thank you. Members, we have heard about the importance of developing and implementing an effective cybersecurity framework for financial and other markets. To further consider these important issues, is there a motion from the body to recommend to the Commission that it re-establish a Subcommittee on Cybersecurity?

(Moved.)

MS. HOUSE: Great. Is there a second?

(Seconded.)

MS. HOUSE: Several seconds. Thank you. It has been moved and properly seconded that the TAC establish a Subcommittee on Cybersecurity. Is there any discussion, any comments on the importance of this subcommittee, and any potential topics that it should be prepared to address, what areas should the subcommittee focus on? Hilary.

MS. ALLEN: Just one quick comment. We've talked a lot about cyberattacks, which are clearly important, but I think resilience also needs to take into account glitches, and some fat finger errors, and things like that. And so I would like to see the subcommittee consider those kind of self-inflicted problems as well.

MS. HOUSE: Definitely noted. Thank you so much, Hilary. Justin, do you have any thoughts?

MR. SLAUGHTER: Yeah. I mean, probably the case is -- question. How many of the panelists here are lawyers?

(Hands raised.)

MR. SLAUGHTER: I am. I think probably most of us are economists. I think probably we need a few people who are hardcore coding experts to participate on the subcommittee as extra members. I'm not. I recognize enough to know that I'm not. That, I feel like, is perhaps the one thing missing from the subcommittee is someone who is up to date on current coding mechanisms, whether it's Python or Rust or whatever, who can speak to the current technologic capabilities.

MS. HOUSE: Appreciate that. Noted. Thank you, Justin. On cybersecurity, I know something that I would propose that the subcommittee consider is the extent to which the financial sector is -- I note since, Todd, you mentioned sharing information, I don't know how robust the -- robust and actionable the information is that's being shared under the ICE AXE, and if they're sharing the right kind of information. Are the indicators of compromise that are being shared good? Is the right kind of information related to risk management? I know that some of that Kevin mentioned is still under way, but the right types of information being shared in both directions, from government to industry, as well as across industry since so many of them are getting attacked with the same vectors. So that's something that I would like to propose.

Any other thoughts or points for discussion from the group, including anyone who's participating virtually, on areas that a Subcommittee on Cybersecurity should potentially consider and examine? Michael, are you coming aboard to make a comment?

(No response.)

MS. HOUSE: All right. Thank you. Then if there's no other comments, then beyond any need for further discussion, we will now take a vote on the motion to reestablish the Subcommittee on Cybersecurity. As a point of order, a simple majority vote of the present TAC members is necessary for the motion to pass.

For those in person, could I please see a show of hands for those voting aye.

(Hands raised.)

MS. HOUSE: Noted. Thank you. Showing the hands of those voting nay on the Subcommittee for Cybersecurity.

(No response.)

MS. HOUSE: No nays. Thank you. For each member participating virtually, please indicate "aye," "nay," or "abstain."

(A chorus of ayes.)

MS. HOUSE: The ayes have it. We will submit the necessary paperwork to the Commission to establish the subcommittee, and we'll be seeking TAC members to serve on the subcommittee.

So we are now ready to explore our third and final topic of the day, Responsible Artificial Intelligence. To begin the discussion, our first presenter will be Alan Mislove, assistant director for data and democracy at the White House Office of Science and Technology Policy. He is presenting on a Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People. Over to you, Alan.

MR. MISLOVE: Awesome. Thank you very much, Carole, for the introduction. As you said, I'm Alan Mislove. I'm the assistant director for data and democracy at the White House Office of Science and Technology Policy. OSTP advises the President and White House senior staff on key issues related to science and technology policy and focuses on coordinating the Federal Government around these policies. Before joining OSTP, I was a professor of computer science at Northeastern University where my research focused on real -- auditing real-world algorithms for issues of bias, discrimination, and privacy leaks. This is the expertise I bring to OSTP where I'm focused on a similar set of issues.

So I'd like to start by thanking the Technology Advisory Committee for the opportunity to speak today, with special gratitude to Commissioner Goldsmith Romero for who -- for her sponsorship of this committee. And I'm thrilled to see that the committee will be focusing on the issue of responsible development and deployment of artificial intelligence.

So today I'm going to be talking about the Blueprint for an AI Bill of Rights, a framework that the White House released last October to help guide the design, development, and deployment of automated systems so that they protect the rights of the American public and reinforce our Nation's highest values. President Biden has pointed to the Blueprint for an AI Bill of Rights as a straightforward set of best practices for both government and industry.

And so here is the problem we at the White House set out to address. Automated systems, often powered by artificial intelligence, now touch nearly every aspect of our daily lives. They have brought many benefits to a range of domains, from cancer detection to agricultural efficiency to helping small business owners cut costs, and we really believe the potential here is extraordinary. But it seems like every day we read another story or another study or hear from another person whose lives have been negatively impacted by these systems. From violating their rights to limiting their access to life-changing opportunities and even to endangering their safety, these systems are having dramatic impacts Americans' lives, often without their knowledge or their consent. And these harms run counter to our core democratic values, values including the fundamental right to privacy, freedom from discrimination, and our basic dignity.

And so President Biden has been clear. We really don't have time to waste in addressing these harms or to protect people's rights and make sure that automated systems work for everyone. To answer this call, and after hearing from hundreds of folks across the United States and beyond, and coordinating with policy experts across the Federal Government, OSTP released the Blueprint for an AI Bill of Rights, which lays out five core protections everyone should be entitled to when it comes to AI and automated systems.

First, safe and effective systems. You should be protected from unsafe or ineffective systems.

Next slide.

Second, algorithmic discrimination protections. You should not face discrimination by algorithms, and systems should be used and designed in an equitable way.

Next slide.

Third, data privacy. You should be protected from abusive data practices via built-in protections, and you should have agency over how your data -- how data about you is used.

Next slide.

Notice and explanation. You should know that an automated system is being used and understand how and why it contributes to the outcomes that affect you.

And next slide.

Fifth, human alternatives consideration and fallback. You should be able to opt out where appropriate and have access to a person who can quickly consider and remedy problems that you encounter.

And so taken together, these five principles outline what we should expect of the systems that are increasingly lee influencing our daily lives. Over the course of building the Blueprint, we listened to people across America, from businesses to engineers, from academics to policymakers, at every level. A key concern we heard was there was a need for resources to help guide the creation of these new protections.

Next slide.

And so in other words, we heard that providing guidance on how to move these principles into practice is as important as the principles themselves, and this problem can be broken down into two questions. First, when, meaning to which systems, should we apply these protections, and second, how should we take these principles into account in the wide variety of automated systems that exist today.

To address the when, the Blueprint for an AI Bill of Rights is focused on protecting people, protecting our civil rights and our democratic values. And so thus, it defines systems in scope based on their impact as opposed to underlying technological choices that they make as such choices can and do change with the speed of technological innovation. And so specifically, the Framework should be applied with respect to all automated systems that have the potential to meaningfully impact individuals or communities' rights, opportunities, or access, defined to include civil rights, civil liberties, and privacy; equal opportunities to education, housing, credit, and other programs; and critical resources or services, such as healthcare or government benefits.

Next slide.

To address the how, the Blueprint for an AI Bill of Rights also includes a technical companion. For each of the five core protections, the technical companion includes examples and concrete steps to build these protections into the technological design process. And so this includes information about, first, why each principle is important, including examples we've seen of problems that happen in practice, and second, it includes what should be expected of automated systems. Taken together, these are the building blocks that are both necessary and achievable to protect the public. And the Blueprint includes examples of how these principles can move into practice, real-life examples of current laws, policies, and best practices that can drive new actions.

I especially hope that you all find the expectations as providing the technical companion useful as these are actionable safeguards that are technologically realizable and necessary. They can essentially be used as a checklist for you or for anybody building, guiding, designing, or overseeing these technologies.

Next slide.

So to give you a sense of what this looks like, I want to dig deeper into one of the principles in the Blueprint, a principle that I think the committee will find particularly relevant, safe and effective systems, which can be summarized as you should be protected from unsafe or ineffective systems. Key aspects of this principle include testing and ongoing monitoring, as well as diverse community consultation, and subsequent reflection on the development and use of the system. Outcomes of these protective measures should include the possibility of not deploying a system or removing a system from use.

Next slide.

The Blueprint addresses why is this principle important. In some cases, models have ended up not working as well in the real world as expected. For example, in addition to underperforming on real-world data, a model developed to predict the likelihood of sepsis in hospitalized patients caused alert fatigue by falsely alerting to the likelihood of sepsis. Making sure that AI systems work with real data is important and so is making sure that these systems work with the people who are expected to be informed by them.

Next slide.

The Blueprint also explores what should be expected of these systems to prevent harm. Here and throughout the Blueprint, we identify concrete steps that can be taken to live up to these principles. This is a checklist that we hope developers and deployers of these systems will use as they implement them. Some of the items that show up on the safe and effective systems include testing and ongoing monitoring. These are basic but important steps we can take to prevent harm, as well as paying attention to the data used is key. When data is created in one context and used in another, it can lead to spreading and scaling of harms.

And so finally -- next slide -- the Blueprint describes how these principles can move into practice. Companies have been instituting many safeguards from internal ethical oversight boards to external audits. In particular, NIST, as you just heard from Kevin, recently released a Risk Management Framework specifically for AI. This RMF emphasizes the socio-technical approach to identifying and managing risks, emphasizing that AI systems do not exist solely in the lab setting, but rather that the safety and efficacy of these systems depends on the societal context they're deployed in and the people with whom they interact.

Next slide.

Notably, the Blueprint pays special attention to sensitive domains where activities being conducted can cause material harms, including significant adverse effects on human rights, such as autonomy and dignity, as well as civil liberties and civil rights. These domains include health, employment, education, criminal justice, and perhaps, most importantly for this committee, personal finance. The Blueprint lays out extra protections that should be expected of systems applied in sensitive domains, including privacy protections for data and provisions to ensure close human oversight and safeguards.

For example, the designers, developers, and deployers of automated systems should consider limited waivers of confidentiality, including those related to trade secrets, where necessary in order to provide meaningful oversight of the systems used in these sensitive domains, incorporating measures to protect intellectual property and trade secrets from unwanted disclosure as appropriate.

Next slide.

So in conclusion, I want to return with to where I started. Automated systems today are influencing almost every aspect of our lives. OSTP has laid out a Blueprint for an AI Bill of Rights as a guide for a society that protects all people from the risks of automated systems and uses technology in ways that reinforce our Nation's highest values. These principles provide guidance whenever automated systems can meaningfully impact the public's rights, opportunities, or access to critical needs or services. Thank you again for the invitation to speak today, and I look forward to your questions.

MS. HOUSE: Thank you, Alan. For our second presentation regarding artificial intelligence issues, we have a presentation from Francesca Rossi, IBM fellow and AI global ethics leader at IBM. She is presenting on The Responsible Development, Deployment, and Use of Artificial Intelligence. Francesca, over to you.

MS. ROSSI: Thank you. Thanks. It's great to be here. Thanks, Commissioner Goldsmith Romero, for inviting me to this session, to this committee, and for learning so much during the whole day. You know, it's really -- it's really great, and I hope that we can learn together from each other.

So my background also is in computer science. I've been a computer science professor for 25 years before joining IBM, and I continue doing research in AI. So I am every day, you know, in contact with the AI research community that has so much, you know, contributed to the recent development, even the one that Commissioner Romero used in the first intervention.

And so in this short presentation, I will give you my idea. I will divide the presentation in two parts. First, I will talk about -- more generally about my vision for AI and AI ethics, and the issues that there are related to AI, and you will see there is a lot of convergence also with the AI Bill of Rights that was just presented. And the second one is I'll give you an example of how IBM, in particular, the company where I work, is handling internally those issues.

So if you go to the next slide.

So this is a very oversimplified history of AI. So I wrote "1956" because that's the time where the term was first used, and usually, you know, AI researchers identified that as the beginning of the adventure of AI. So I'm oversimplifying a lot, and I'm giving you this idea of the history because some of the issues will be also related to the different kinds of techniques that are used in AI.

So first, there was the so-called symbolic, or knowledge-based, or logic-based artificial intelligence where, basically, people were writing down algorithms that could solve intelligently a problem, and then these algorithms were coded into the machines. So then -- from then on, the machine was able to solve the problems intelligent, but then intelligence were given by the people quoting those algorithms, right, so telling machines what to do one step after the other on how to solve a problem.

Then in the 80s, there were the introduction of a completely different way of telling machines of how to solve a problem, and this different way was based on data and techniques called machine learning. And the idea was to tell the machine not the steps to solve the problem because in some situations, you don't have that luxury to be able to tell the machine all these steps. Like, for example when you try to recognize a face, or an object, or a cat, or a dog in an image, you cannot tell the machine the steps and be sure that at the end, they will be able to recognize correctly. So you have to give a lot of examples of problems and their solutions and then let the machine learn from these examples to solve the problems also for images, for example, that it doesn't -- it doesn't see in the example. So ability to learn from data without being explicitly told all the steps to solve the problem.

These techniques were around in the research community since the 80s, but they were used and practically used only much later because they need a lot of data and a lot of computing power to work well, and we didn't have data nor computing power that was enough in the 80s. It were only later that we started uploading so much data on the web, and we had the internet, and so on. Then these machine learning techniques, and notice that the other techniques continue. They're not shut down and moved to machine learning only. Then the deep learning techniques were -- came about around 2010, which were based on these deep neural nets with several layers, like the picture that you see there, that helped optimize and scale these machine learning techniques.

And then lately, the generative AI, which is still based on data -- of learning from data, but not only can interpret well images, text, and understand what is in those images, what is in the text, and so on, but can also generate images, text, videos, and so on. So that's why it's called generative AI, and ChatGPT is one example of a generative AI technique. In fact, "GPT" means generative pre-trained transformer, which is one specific technique that is used to -- for this generative AI. So and that, for example, is -- the image that I put there is an image from Dali that is another generative AI example the produces images from text.

Okay. If you go to the next slide.

So AI is used -- we don't even know how -- that we use it all the time, you know, all the time. In everything we do online, we use AI. It supports many of our activities. But, most importantly, probably for this session here, is that it is used also in many high-stake decision-making applications, like financial institutions, not relevant here, but also H.R., employment, admission to schools, healthcare, all the workflows of the enterprise. So really, AI plays the role of supporting a lot of decision environment where the stake is very high. So that's why we need to be careful about the issues related to this technology.

And if you go to the next slide.

And I'll give you some -- a very incomplete list of some of the ethics issues that many, you know, are talking about. So first of all -- and I will tie this issue to some of the characteristics of the AI techniques. So the first one is data privacy and governance, and why is that a central issue? Because as I told you, machine learning techniques need a lot of data to work well. Generative AI needs even more data to work well, so AI data privacy, and governance, and sharing, and collection are really central issues.

The second one is fairness. So the issue that AI can make or recommend decisions, and, of course, just like we don't want this decision to be discriminatory from a human being, also we don't want them to be discriminatory made by AI or by human being recommended by AI. And the reason why AI can make discriminatory decision is that it's trained on data that can contain some bias because we generate the data, we collect the data, and there may be some correlations from variables in the data that can pick -- be picked up by the machine learning algorithm that then, when it makes a decision, can use that correlation to generate discriminatory decision. So that's something to be really careful about.

Another one is that AI should be used in a way that is not creating gaps, so it should be inclusive. So that's related, for example, to the opt-out, to the ability to say, you know, I want to use a person and not an AI. So inclusivity and fairness are related, but they are two different things.

Then there is explainability and transparency. So explainability is a property that should be of the technology. So the technology should generate an output, and then it should be able to explain you why it generated that output. If my loan application is rejected, I want to know why it has been rejected, so the technology has to have that property. Transparency instead is more a property of those building the technology. Those teams, companies that build the technology, they need to be transparent so that whoever is using that technology in another company, for example, can have a more informed use.

Accountability is a word that has been used a lot today, and, of course, it is important here as well for AI because AI, especially machine learning, is based on statistics and probability, so it always has a small percentage of errors. And so it's important to understand who is accountable when things are not done in the right way. Social impact. AI, as you have seen, it devolves very, very rapidly and is very pervasive, so it generates a very fast affirmation -- jobs in society -- and we need to understand what to do about it.

The second slide about the ethics issues is about human and moral agency. AI can profile people if it has a lot of data about the person and even manipulate our preferences, and then there are some issues you see related to -- especially to generative AI. So AI can generate content that seems very plausible because it has a very high-quality fluency of the content, for example, the text that is generated, but the content may be false, and not everybody goes and checks that content.

So that -- the issue of spreading possible missing information, as well as value alignment because we have seen that the text that is generated can be harmful, can be toxic text, can be inappropriate, offensive, racist, and so on. And this is not because AI has some malignant idea of generating that context, but because generative AI has been trained to do one thing only, for example, large language models, to generate the most probable next word out of the previous, like, 300 words. And so it's not that it's voluntarily or own purpose is lying or generating that context. It's that we need to understand how to embed the values in that AI if we want to be aligned to our values.

As well as some issues about environmental impact, generative AI especially needs a lot of energy, a lot of computing power, and a lot of data to be trained and also deployed. Also power imbalance, so the amount of data and computing power that is needed is really so much that not everybody can afford to build such a model.

Next slide.

So that's where AI ethics come about, tries to address all these issues, taking the best of the technology but mitigating those issues. And it does so in a very multidisciplinary way where AI experts get together with all the other experts in other fields so that understands the impact of the technology on society and many other aspects. And that's why the solutions that AI and this field put together are -- some are technical, but many are social-technical, and it's the puzzle of solution that I will talk to you about.

Next slide.

Over the last, let's say, almost 10 years, we have seen three phases in AI. In the first phase, few years, people were just trying to identify these issues. They were seeing things that were not going well, but the issues were not really completely identified at the beginning. Then there was a second phase of principles. Everybody wrote principles, you know, national government agencies, and companies, academia, civil society organizations. Everybody, and multi-stakeholder organization, everybody wrote principles. And now we are in the practice phase, and this is reflected also in the principles intervention. Yes, principles, but -- principles and rights but also how to translate them into practice.

And this is where we are now. Everybody is in the practice phase: regulations, standards, corporate, internal directives, processes, auditing, certification, and so on. But not only AI has evolved in those three phases over the last, let's say, 10 years, but also, as you have seen, AI has evolved as well. So generative AI was not in the picture when there was the awareness phase and the issues were identified. So that's why now, yes, we are in the practice phase for those issues, but we also have to be fast because there are new issues that are being introduced, so new principles and new practice. We have to be much faster now rather than taking those seven or eight years.

Next slide.

Okay. We said already social-technical issues, they need social-technical solutions. Yes, tools, but also education policies, multi-stakeholder consultation, and many others.

Next slide.

So here I give you an example of the fact that every societal actor has to be involved in addressing those issues. Research communities, and there is a lot of AI research around how to build technical solutions to fairness, explainability, all the issues that I mention. A lot of AI companies building or using, deploying AI, governance internal processes, tools, risk assessment, and so on.

Standard bodies, and here I got -- I just give one example by IEEE, but there are many standard bodies that really focus on a AI -- standards for AI issues. Educational institution, and here, I just gave a list of courses all over the world, educational institutions, universities are really trying to add courses about the ethics of AI given together to science students, as well as governments, and, of course, AI Bill of Rights is one. AI in Canada, a European AI Act, which is still very actively discussed at this point and probably will soon be approved, and many others. But these are just some actors, and they don't work in isolation. They work together but also together with civil society organizations, media, society at large, so everybody needs to be involved.

Next slide.

So let's focus a little bit about what a company can do to address these issues, and the question is -- that any company would ask is why should I invest money, and time, and effort, and people in activities around the AI ethics. And many companies would say I'll wait for the regulation. When the regulation will come, I will comply, and this is a very shortsighted, in my view, approach because, yes, of course you have to comply when the regulation will come, but you -- first you have to have some company values that you are going to be aligned with independently of the regulation because, otherwise, the -- your reputation of the company, the impact that you will have on society -- on society for your company, the trust of your clients will be impacted independently of whether there is or there is not regulation. And also, you will get even more business opportunities if you anticipate regulation. And on top of that, there is value in working with regulators, in helping them define regulations in the most informed way, also from a technical point of view.

Next slide.

So let me give a very short overview of what we do at IBM. As everybody else, we started already in 2017 with our principles, very high-level principles. First, we think that AI should augment human intelligence and not replacing it. Okay. Second, data that we collect from our clients, one client, we don't reuse it from another client. They belong to creators. A third one, the focus on transparency and explainability, and we call this principle "principle for trust and transparency" because trust is a central focus of our framework.

Now, you see these three principles are very clear, very nice. Who could disagree with this principle? But they are not useful for developers, for our consulting division. They are not useful. We have to go down -- much more down into concrete action. So that's why --

If you go to the next slide --

-- from the principle, we said, okay, how do we structure all our activities around AI ethics, and we structured them around these five pillars of what we call trustworthy AI. And you see, you recognize here many of the central issues that I told you about as well as the rights that were described earlier: explainability, fairness, robustness, transparency, and privacy. So we want to deliver technology where these five pillars are addressed in the best way.

If you go to the next slide.

Okay. So it's very natural for a tech company to do what? To think that technical -- technological issues can be solved with more technology, and in some sense, this is one important piece of the puzzle. And this is a list of tools -- many of them are open source, but also, some our proprietary -- that IBM put together, each one devoted to mitigate one of the issues -- fairness, explainability, and so on. So you see AI explainability 360, AI fairness 360, and so on, so very, very useful tools. But soon after releasing these tools, using them, and we still use them and our clients use them, we realized that the tools are just one piece of this big part of the test to be put together, which includes many other dimensions.

If you go to the next slide, you'll see that there is a very, very powerful governance structure.

We have an internal AI ethics board that is co-chaired by myself and by the chief privacy officer, that is positioned well in the governance of the company, and that's the decision power. So it's not an advisory board. It's a decision power to decide how our developers are building AI systems and whether an AI solution is delivered out of the company to a client or not. So we have risk assessment for each solutions. We have playbooks for our developers, so, to tell them how to use the tools and how to consult with all the stakeholders to ensure that the right properties are there.

The AI ethics board communicates as members of the board, come from all the business units, not just the technical ones. And it communicates directly with every business unit, not just with the member but also with we call the focal points that are people in the business units that achieve this by directional communication between the business units and the board. The board makes decisions. Those that are relevant to the business units are brought down by the focal point, and the focal point brings up challenges, feedback, grassroots initiatives, and so on. So very important to have is a very powerful governance structure, not just an advisory board.

Next slide.

These are some of the activities that we do and this board supervised. So as I said, tech ethics by design is telling our -- the playbook for our developers, how to integrate AI ethics in the development pipeline. The use case reviews, this is how we assess that the deal should be signed or not, whether it's aligned to our principles or not. Collaborate with policymakers, of course. Educational modules for developers, very deep, but also for all the other IBM networks. A workstream of foundation models, very recent of course. Connection with Eurotech, and also the focus on really what it means to augment human intelligence rather than replacing it.

Next slide is about the value of multi-stakeholder collaboration. So you will -- I told you, the AI ethics board is an internal board, does not have external people because we think that the best way to collaborate with the others is by collaborating in partnership. So there partner -- we are founding members of the Partnership on AI. We work with the World Economic Forum, with academia, like MIT or Notre Dame University, with IEEE, with the U.N. agency, like ITU, and also, I was a member of the European Commission Expert Group on AI that define the ethics guidelines for trustworthy AI in Europe. So a lot of multi-stakeholder collaboration.

Next slide.

Okay. This is my final slide. So some lessons learned in trying to operationalize AI ethics principle. So as a first thing is that it needs to have -- a company needs to have a company-wide approach where all the divisions are involved with their different roles, not an AI ethics team that works only the team that then tries to connect with the different divisions. It has to be a company-wide approach.

A governance body, as I told you, with the power to make decisions. If the -- if the AI ethics board says this deal cannot be signed, there is no way to make it align to principles, the deal is not signed. Full operationalization of the principles, which means, yes, the principle are fine, but then you have to go down to very detailed and concrete actions. Tools are very useful in a tech company and to address technical issues, but also all the other pieces -- processes, education, risk assessment, governance, and so on.

Regulations are, of course, important, but AI ethics is beyond the compliance regulation if you want to be really looking at the long-term sustainability of using certain technology. And the value of multi-stakeholder partnership to learn and also to bring experiences, challenges, and learn as well. The last one is that as AI evolves so rapidly, all these approaches, even the one that we put together, are evolving. So, for example, the five pillars that you saw there, they are evolving, and they now add more pillars that have to do with issues that are present in current AI and maybe they were not present even a year ago because there was no generative AI that was available for everybody to use.

If you go to the last slide, you have some information that you want to look at more at the -- our approach to deal with AI ethics, that QR code takes you to our website. And the other one is a World Economic Forum white paper that was written last year, 2021 actually, to describe to, in a succinct way, our approach. So if you go to our website, you'll find a lot of information in that website -- in that white paper. You'll find all the things that I told you about in a very succinct way. Thank you.

MS. HOUSE: Thank you, Francesca. For our third presentation regarding artificial intelligence issues, we have a presentation from Tim Gallagher, managing director at Kroll, regarding the Emerging Threat of AI-Enabled Cyber Attacks. Tim, over to you.

MR. GALLAGHER: Great. Well, thank you very much for the opportunity to present today, and I appreciate the opportunity to be on this committee which is addressing such important work as far as threats to our Nation's financial infrastructure. As you heard, I'm Tim Gallagher. I'm a managing director at the firm, Kroll. I'll give you a quick background or

Kroll and then myself, which I'll provide some context for my remarks.

Next slide.

Kroll has been around for about 50 years. It's a risk management firm based in New York City. We protect -- we work with clients to protect their people, their property, and their reputation. And as everyone in this room knows, cyber impacts every last bit of that. Worldwide, have about -- over 600 practitioners working on cyber. We do about 3,000 incident responses a year.

Next slide.

My background, I came to Kroll five years ago after a 22-year career in the FBI. I held senior positions in the Cyber Division in FBI as well as the Financial Crimes Section. My practice here at Kroll is where cyber meets fraud, which is pretty much everywhere as well.

Thank you.

So in framing this up today, obviously this headline you see here was taken from The Wall Street Journal five years ago. However, it could have been five days ago. Every time you pick up the paper, there's something else about artificial intelligence, whether that's platforms, such as ChatGPT or, you know, the latest competitor to it -- I guess it was Bard this morning, which was announced that we're seeing out there, as I said today.

The next slide.

I'd like the frame this up by looking at -- looking at some of the numbers here. ChatGPT, in the last six -- in a six-week period went from one million users to 100 million users, so, like, what does that tell us? You got to look at that as a -- you know, with my FBI background and with my cybersecurity background, what do I do right now? At Kroll, you have to think like a large percentage of those users are people who are exercising their, you know, what they consider their God-given right to defraud the American public, so they're using this as a new tool to go out there and try and commit fraud schemes.

About a year ago, the FBI put an alert out on foreign influence operations. You know, I'm sorry. Spoiler alert here. It's going to say that cybersecurity schemes with AI will be focused on phishing lures and misinformation campaigns, and that's exactly what we've been seeing. The FBI Internet Crime Report, which came out a couple of weeks ago, shows $10 billion in fraud that comes back to online scams. What portion of that is attributable to a AI? We don't know. It's too early to tell. However, you can be sure that that number is going to go up because the barrier to entry is now lower where fraudsters can just utilize this new -- these new tools out there, who may not have been able to get in before, as well as threat actors who are actually good at what they do or able to use it to enhance their schemes.

Partnerships. That's my last bullet there. Law enforcement, which is where I came from, and regulatory agencies, like the CFTC, they can provide intel. They can get information out there to the general public, but the companies out there, you know, you have to do the mitigation yourself. You have to take that information and work -- you know, work with your internal components, work with each other, and work with other governmental agencies to mitigate the threat.

Next slide.

So here's what we're seeing at Kroll. As I said before, we did over -- do over 3,000 incident responses per year, and we're pulling information in, and so I just want to give you an overview of what we're seeing so far.

Initially from the threat intel groups, so these are the folks that we have out there who lurk in the dark web, you know, who lurk in these -- in these hacker forums and see what they're -- you see what they're pulling in. And right now, what we're seeing is discussions about ChatGPT using it to create malware, and -- but it's not really happening yet in a way other than what my folks are telling me, that they consider to be publicity stunts. You know, maybe someone wins a hackathon by using ChatGPT, but the software itself, from what they're seeing and the chats that they're seeing, is being utilized more to generate verbal-type of interchanges. Kind of like, you know, Commissioner Goldsmith Romero, your opening remarks, right, that were partially written with ChatGPT. As far as actually writing malware, there are other programs out there which have been around longer which do a better job. So right now, they're not actually seeing it.

However, interestingly enough, they are seeing in the chat groups introduction of code that will get around the user agreements that go along with ChatGPT, i.e., trying to get it to be utilized in countries that are on the blacklist for ChatGPT. So there's code out there that's trying -- that hackers are talking about introducing to jailbreak ChatGPT, and so they're able to use it for nefarious purposes that it was not intended before.

Okay. Next up, our malware analysts. These are the individuals who are, you know, they're always pictured being people who are sitting there with their hoodies, you know, behind the keyboard in every one of these memes that you see. But like the ones who work for us at Kroll, they really do look like that, so when I have my meetings with them once a week, they get up on the screen and they have the -- you know, they have their headphones that are -- you know, they're gamer headphones, and they have their hoods up. The only difference is that they used to work for the NSA or they used to work for the Secret Service, or they worked for GCHQ. You know, they're hackers, and we hand the malware and we say find the evil, and they -- and they get all excited about it, you know? We give them here's what we're -- here's what we pulled in from one of our recent engagements, you know, what can you tell us about it.

And what they're saying is, like, there's going to be pieces in there that are AI generated that we're seeing right now, that maybe even pulled out of -- you know, utilizing ChatGPT, but it's what they call the cut and paste portion of it. It's not the real -- the real hacker and real hacker tools that are being utilized to penetrate networks. There's some of it, but like I said, it's not the real -- the real cutting-edge tools. It's just going to be the cut and paste.

Also, business -- go down to the next -- the next box I have up there, business email compromise, phishing attacks.

My last assignment in the FBI, I was special-agent-in-charge of the -- of the Newark office. I was in charge of operations of the whole State of New Jersey. I used to love to tell people that every crime, no matter where it starts or where it finishes, at some point, it's going to go through New Jersey, so it was, you know, a great background to have.

I say the same thing about phishing attacks, like, and business email compromise. Every cyberattack on some level has a phishing attack in it. By the end of last year, 77 percent of what we saw coming in to Kroll, the initial access, the initial infection was done with a phishing attack. We talked before about, you know, hey, don't click on that. That's further broken down: 69 percent is going to be clicking on a link that you shouldn't have clicked on, and the other eight or nine percent is going to be an attachment like we saw in the -- in the blockchain attack where it was a PDF that was opened up.

So as you can see, in summary, about 77 percent of what we're seeing involves malware that was introduced through a phishing attack. Fifteen percent involved vulnerabilities. You know, someone spoke before about zero days, so that's the number that we're seeing as far as the initial infection coming in through zero days, about 15 percent. So still, you know, the vast majority of cyberattacks on some level are going to involve a phishing attack for initial access to the system.

Business email compromise. FBI figures: $2.7 billion last year. Thirty-three percent of what we're seeing are business email compromises. You know, why am I focused on that? Because that's where we're most likely to see more AI and more ChatGPT attacks right now. As we -- as we know, a lot of these attacks involve non-native speakers, so this will give them the opportunity to write better fishing lures to get themselves into the system, to get -- to have you more likely to click on something that you maybe should not have clicked on.

And then once you're in the system, as you know, the business email compromise is going to involve some banter back and forth to try and get you to wire money

-- the victim to wire money to some place they should not, and that's where the chat -- the chat function, the bots will -- are being utilized. They will be utilized more as a way to get you across the line and wire money where you where you should not have.

Threats to the markets. Obviously, everyone in this room knows how critical the -- to the investing public is confidence in operations in our Nation's financial markets. You heard from our colleague from NIST before, you know, their job is to cultivate trust in technology. We need to have trust in markets for folks who invest -- to invest in the market. When I was in the FBI, we had the Fair Markets Initiative, and that was, you know, to keep the public's confidence in the market.

The pump-and-dump scheme, which is something that we've all seen here, it's evolving where to the point where these AI functions make it easier, right? Like, breaking down the elements of that crime, number one, you're recruiting people in to invest, and that's where the, you know, the AI and ChatGPT function can help out by going out there and finding folks who, you know, based upon what they have out there in their profiles, would be likely to invest, or putting out synthetic profiles where they're hyping certain stocks as a way to get people in or certain commodities. And then, of course, the hype function is also extremely important, and that's where writing the fake news or the -- or the -- or the releases that's going to, you know, pump that stock up is going to be and we're seeing is being affected by AI as a way to write something that's going to get you to invest and pump that stock up.

And then lastly, the public/private partnerships. As I said before, the -- we need to -- all need to work together on this. I've seen it, as I said, from the -- from the -- from the government side and now from the private sector side, and the amount of sharing that we're seeing out there, I know someone brought this up before, that this needs to be an aspect of what we look at in the -- in the Technical Advisory Committee. You know, I couldn't agree more on that.

It's absolutely amazing, having been on the government side where it was all a one-way street, where we would take information in from the private sector. We would not really push information out, but, of course, that was 20 years ago. Now the flow of information is absolutely amazing. I get bulletins from CISA that puts out their TTPs of what the threat actors are doing right now, pretty much real time. My folks are in slack channels with CISA where they're getting information about what the latest attack schemes are.

The FBI has keys to unlock some of these ransomware groups that are out there, and they share them with the private sector. So, you know, the FBI wants you to report ransomware attacks because a lot of folks don't, and that's the carrot they have. Hey, look, if you report, we may have the keys, you know. We could help you -- we can help you unlock that. You know, the CFTC putting out bulletins, like, here's what we're seeing scam-wise, and here's what you should be on the lookout for.

So it's an exciting time for public/private partnerships, and I'm happy to be leading -- I'm sorry

-- being a part of that initiative being led by Commissioner Goldsmith Romero. So thank you for your time, and I'll join any questions with everyone else here.

MS. HOUSE: Thank you, Tim, and great shout-out to the FBI's hive actions. So at this time, I would like to open the floor to questions and comments from the TAC members related to responsible and ethical use of AI.

MR. CUTINHO: This is a question for Francesca. You know, how -- I mean, IBM is a global firm. It's not -- it's not a question directed at IBM. It's a question directed at any firm -- private firm that is getting into this business. So, and you do business in multiple jurisdictions. It's admirable that you're taking a lot of steps on the ethical side, but my question is how does IBM protect itself if one of the institutions in any of the countries you operate in or an international institution, coerces you to manipulate using your technology. How would you protect individuals from that?

MS. ROSSI: So how would that happen? I mean, any --

MR. CUTINHO: I mean, how do you protect you -- I mean, how do you protect us? We are vulnerable, right, so because -- and in many ways you are vulnerable because you're regulated in many of these countries.

MS. ROSSI: Yeah.

MR. CUTINHO: And there are many international institutions that have great influence over you. So if you're coerced to use your technology to manipulate populace --

MS. ROSSI: Yeah. Well --

MR. CUTINHO: -- how would you -- what would your reaction be? I mean, that is the most difficult question -- ethical question, isn't it?

MS. ROSSI: It is an ethical question, but the framework that I showed you that we use in these centralized governance that we have -- the AI ethics board -- is the same one, same framework, same thresholds, same all over the world. So it's not for U.S. and then we have another one, maybe it's a bit more relaxed for another region of the world, or another one more rigid for another. So it's the same one for all our deals that we signed all over the world. So whether it's -- you know, everywhere.

So I haven't seen so far, since the board was in place and there's been several years already, any situation in which what you mentioned, you know, happened or was, you know, going to happen if we didn't take measure. We evaluate very -- many, many different use cases coming from all over the world, the teams that want -- that want to sign a deal with a client, and the -- this deal raises ethical issues.

So it comes -- the team comes to the board and discusses with the board, and the board helps understand the team how to make the deal align to our principles, for example, by requesting more tests, for example, for bias and, you know, that -- and seeing the results of the test, or by adding some specific terms and conditions on the contractual agreement. That's another thing that we impose many times, for example, related to inclusion. Like, I remember when IBM was working on the digital health pass for New York City, we imposed a paper version because we didn't want to be known inclusive in the deployment of that technical solution.

So that's what we do, but the discussions are the same. The only thing that changes is the team because the team can be local, the team that is going to sign the deal, but the decisions are based on the same framework and same thresholds all over the world.

MS. HOUSE: Thank you, Francesca. I appreciate the complexity of the issue and the answers, like, looking at whether its governance policy, tech controls to prevent exploitation, deal with cross-border issues, whether the data is across borders or the application. So a really interesting issue coming up. So I've got -- I see several other flags, so I'll just ask folks to keep their interventions at under a minute if that's all right. Hilary, you're up.

MS. ALLEN: So just quickly, I'm really interested in the use of AI in the context of risk management, which, I think, would be pertinent to the CFTC's ambit. So I think something to think about is the idea of when we are in a Big Data situation and when we're not, data quality is key to this stuff working. And if we think about it, there really is, in some respects, only one market history. So I think it's just something to think about when we're actually in a Big Data situation or not.

MS. HOUSE: Thank you very much, Hilary. Dan Guido.

MR. GUIDO: One thing that I'd like to point out is that with a lot of the generative AI systems that have popped up, they operate in very unconstrained matters. Like, you can -- you can ask ChatGPT to be a lawyer, a doctor. You can ask it to review a cancer scan of you. You know, you can do a million different things with it, so it's very hard to figure out is this thing fit for purpose. What is it supposed to do? If there's no, like, use case to find or specifications that it operates against, it's very hard for an outsider to evaluate whether it actually meets the goals that we've set out for it.

So from an assessment perspective, one technique that we find very valuable at Trail of Bits is the use of something that the self-driving car folks came up with called an operational design domain, ODD. And with an ODD, you're able to define what is this system intended to be used for and then measure its performance against those goals. So while the risk management framework from NIST is an excellent contribution to the space and the Bill of Rights is an excellent contribution to the space, from a risk assessment perspective rather than a risk management perspective, I think ODDs are a positive contribution.

MS. HOUSE: Thank you.

MS. ROSSI: Can I -- can I respond?

MS. HOUSE: Sure.

MS. ROSSI: I mean, it's not a question. I understand. You're right that there are two ways to use, for example, a large language model. One way is to use it is -- in this open-ended domain situation, like ChatGPT. You can ask any question about anything. And another way is to use it as a foundation for building very specific AI systems for a specific task. And then in that -- in that case, which is what IBM is doing with -- that's why we call them foundation models because we see them as the foundation for building a specific AI system.

And then you can add to the vast amount of data that the large language model is trained on. You can add additional data coming from the specific tasks, and the client, and the company, which is curated data. So that allows you to build a specific system for a specific task that can also mitigate some of the issues that have to do with misinformation of affect duality that happens much more in an open domain.

MS. HOUSE: Thank you. I'm going to turn next to Nicol. I think you have a question --

MS. LEE: I do, and I apologize for not being there and very honored to be on the Commission. I was in between meetings. I've been listening for the last few hours. My question for the panelists just as this Commission begins to think about AI and, particularly, generative AI, what should we be concerned about because we do know that generative AI is obviously sort of this next deeper, potentially more efficient type of artificial intelligence when it comes to its ability and its cognition. But at the same time, in a very fluid, if somewhat vulnerable economic system and financial markets system, that has a potential to be -

- to cause unintended consequences.

And the work I do at Brookings has a lot to do with unintended consequences for vulnerable populations, but I think in all the presentations we've had today, there's also the potential for generative AI to be weaponized in ways that traditional measures of cybersecurity and other gap stops, you know, may not be able to manage. So I'm just curious as we think about the formation of a subcommittee around this, and even bring it back to the White House and what you're trying to do to put together a Bill of Rights that sort of mitigate some of these risks that may occur on the socioeconomic side or social side, I'm just curious from any of the panelists, the types of things you think we should be anticipating in this space with the emergence of generative AI.

MR. MISLOVE: Great. So, Nicol, I think that's a -- that's a great question. At the White House, we really developed the Bill of Rights as a resource to answer those kinds of questions for anybody developing, deploying, using, trying to regulate AI.

And so the latest wave of generative AI as sort of the new thing we need to be thinking about.

And the way we wrote the Bill of Rights, we talked about the impacts of these systems and not how the systems are implemented. So the AI Bill of Rights was released in October. It was ChatGPT became sort of a thing. And so the AI Bill of Rights is really focused on not technical details but sort of how the systems are used and how they could be implemented.

So the first place I would start with is looking at the principles there and sort of how to move those principles into practice as a way of framing what are the risks and then how could we move towards mitigating the risks of these kinds of systems. For example, the things about harmful output and misinformation would be violations of safe and effective systems, and there's also issues of privacy that some of the other panelists brought up that would be violations of privacy protections.

MS. HOUSE: Thank you, Alan, and thank you, Nicol, for the question. Joe, I see you have a comment or a question.

MR. SALUZZI: I've got a question. Thank you first. I mean, I come from the equities market, so I've learned a lot today. Thank you very much. But my question actually is for Tim. I study a lot of market structure in the equity market, and the problem that we had two years ago was GameStop, right, the right the GameStop phenomenon back then. We had three congressional hearings. We've had numerous interviews and so on, and no one quite figured out exactly what happened. And nd a lot of fingers were pointed at, you know, market structure, but certainly at Reddit and what was going on there. So do you think that this, now with ChatGPT and all the AI, could that just get another GameStop? Could we have another problem like that and even quicker this time?

MR. GALLAGHER: Well, that's really good question. We haven't figured out what happened the last time yet, and, you know, are we poised to handle the next one. You know, all I can say is, like, the good guys have the same tools as the bad guys, right? So we have to be using these tools to try and get ahead of it, and potentially put some stops in place, and utilize that to try and -- try and mitigate it before the next one happens. That's, you know -- that's all I can say to that, yeah.

MS. HOUSE: Thank you. All right. Then Members, we have heard about the significance of emerging technologies, such as artificial intelligence. We've also heard and discussed earlier evolving technologies, like cloud solutions. To further consider these important issues, is there a motion from the body to recommend to the Commission that it establish a Subcommittee on Emerging and Evolving Technologies?

(Moved.)

MS. HOUSE: Is there a second?

(Seconded.)

MS. HOUSE: I think so. Thank you. It has been moved and properly seconded that the TAC establish a Subcommittee on Emerging and Evolving Technologies. Is there any discussion, especially any comments on the importance of the committee, the nature of the work that we think that the subcommittee could especially focus on and contribute to the Commission?

We'd love to hear views from anyone, just brief views.

(No response.)

MS. HOUSE: I know one thing I would -- that I'm excited very much about the prospect of the subcommittee, if the -- if the committee votes to approve establishing this or to recommend to the Commission to establish it, looking at the -- any specific applications related to supervisors that we think, out of all the principles that have been laid out, what are the ones that maybe are the stickiest or the ones that are needed earlier on to enable oversight of some of the other principles, things like explainability and transparency, that those be needed in order for our supervisors to properly regulate and oversee whether or not the other principles are, in fact, being followed inside of those -- inside of those AI solutions.

So that's something I think looking at whether there's -- whether there's any lessons, or applications, or principles that should be focused on first, especially related to the supervisory function that would be relevant to all regulators and any others trying to make sure that AI is being operated and designed responsibly. Any thoughts from others about the areas of focus for the subcommittee to look at? Yes?

MR. SIRER: So I think a lot of the ethical dilemmas here come from autonomous decision-making systems. One thing that I'm curious about is the emergence of decentralized autonomous organizations and the ethical dilemmas they pose because they are able to do things and take actions that regular organizations with centralized, you know, means of control cannot do. So that's an exciting, at least from my perspective, an exciting area.

MS. HOUSE: I think that's a really interesting one that could point to, like, if there's specific innovations in finance generally, including the use of Dows and others using AI, the unique sensitivities of privacy and governance needed for Dows, although I realized Dows can operate non-financial systems as well. But I think that's really interesting autonomous decision-making point.

MR. SIRER: Absolutely. Just so that everybody is on the same page, one of the exciting or one of the -- one of the interesting things that came up when the DAO debacle happened on Ethereum was that this thing, the DAO, could fund activities that you could not normally get funding for, such as somebody coming before it and borrowing funds to build a submarine, to ferry drugs from one part of the world to another. And these are now becoming -- it's at least theoretically conceivable.

MS. HOUSE: Thank you, Grun. I appreciate that. Then, Corey, if you'll make our closing comment before our closing vote.

MR. THEN: Sure. Thank you. I think it might be interesting to have some sort of gap analysis with existing consumer protection laws. So, you know, we have Fair Housing Act, TELAF, FCRA, all these things, and outside of just the ethics component, which is on top of laws, to figure out, okay, do any of these sort of like existing rules sort of ameliorate concerns. And if not, and they're not covering where the technology is moving, what might we need in addition to them?

MS. HOUSE: Thank you, Corey. If there's no further discussion, we will now take a vote on the motion to establish the Subcommittee on Emerging and Evolving Technologies. As a point of order, a simple majority vote of the present TAC members is necessary for the motion to pass.

For those in person, could I please see a show of hands for those voting aye.

(Hands raised.)

MS. HOUSE: Thank you. A show of hands for those voting nay.

(No response.)

MS. HOUSE: For each member participating virtually, please indicate "aye," "nay," or "abstain."

(A chorus of ayes.)

MS. HOUSE: The ayes have it. We will submit the necessary paperwork to the Commission to establish the subcommittee, and we'll be seeking TAC members to serve on the subcommittee.

MR. BIAGIOLI: Well, thank everyone so much for coming. This has been a very productive and illuminating inaugural meeting. I know I feel that way. I think -- I won't speak for others, but I'm sure others feel -- Commissioner Goldsmith Romero, do you have any closing remarks?

COMMISSIONER GOLDSMITH ROMERO: I do, but I first want to hear closing remarks from our chair, Chair House.

MS. HOUSE: Thank you, Commissioner. This his has been a wonderful day. I'd just continue to reinforce how honored I am to serve amongst all of you. I wanted to highlight as we stand up these subcommittees, or if the Commission determines to move forward with the recommendation from the committee today, that I really look forward to vigorous participation from everyone and debate. As we've seen earlier, there are certainly some points of disagreement, differing perspectives and expertise, wisdom and experience across this -- across this entire committee, so I believe that the Commission will benefit most from especially seeing the areas where we disagree actually. So I'm really looking forward to diving into that with all of you. This was an incredible foundation. Thanks very much to all the presenters as well and to the Commissioner.

COMMISSIONER GOLDSMITH ROMERO: Thank you. This has just been wonderful. I mean, I'm sure you can see the very serious implications that come from technology just in what we've discussed today, and there are many topics we couldn't get to today, but I knew that this team would want to go deep. I know on any one of those topics, we could have just had one presentation and probably talked for hours and hours and gotten a lot more viewpoints. So to narrow it down to three was a little tough.

So I think one of the -- one of the best things about what I saw today is how willing all of you are to share your expertise, to share your perspective, to share your views. I would -- I would suggest that you continue to do that amongst each other you. You all have each other's emails. You all can reach out. You all have context now. These are -- there's a lot of topics here that are worthy of discussion, either that we discussed today or that we should discuss, and I'd certainly like to hear your views.

I know Chair House and Vice Chair Redbord would like to hear your views as well. And so, you know, please use this as a group to really be thinking through these things, to work things out, to give good ideas. Ultimately, all of this is incredible advice for the Commission. There are a lot of people who watch this live, and there are a lot of people who will watch this later, and so you never know the influence that you're going to have on policy making.

So I'm very, very grateful for the expertise. I'm grateful that everyone is willing to share even their different views, and respect each other, and really take this seriously. You know, we've brought together people who are very, very serious in terms of thinking about this, and you think very deeply on these issues, and we want to make sure that all of that is brought in. And so I'm grateful for your service.

I will just say this is public service. None of you were paid to be here, and, you know, it's not like we can do anything other than to offer you our gratitude and to recognize that public service, and so thank you. I look forward to the continued work. I'm already very excited. This was a very energizing day for me, and I'm grateful to have all of you serve here, and it's an honor to sponsor this group. Thank you.

MR. BIAGIOLI: With that, thanks so much to everyone once again for attending our first TAC meeting. The meeting is adjourned. Have a great evening.

(Whereupon, at 4:55 p.m., the meeting was adjourned.)