QFRS
Senate Judiciary

September 21, 2022

Peiter “Mudge” Zatko
Independent Security Consultant
New York Metropolitan Area

Dear Mr. Zatko:
Thank you for your testimony at the Senate Committee on the Judiciary hearing entitled “Data Security at Risk: Testimony from a Twitter Whistleblower” on Tuesday, September 13, 2022. Attached are written questions from Committee members for your review and response. We look forward to including your answers, along with your hearing testimony, in the formal Committee record.
To complete a timely and accurate hearing record, please submit an electronic version of your responses to record@judiciary-dem.senate.gov no later than 5:00 p.m. on Wednesday, October 5, 2022.
In the case that circumstances make it impossible to comply with the time period provided, witnesses may request an extension to the above email address. Any additional questions, comments, or concerns may also be directed to this email.

Sincerely,

Richard J. Durbin
Chair
—--

Senator Chuck Grassley, Ranking Member
Questions for the Record
Mr. Peiter “Mudge” Zatko
Independent Security Consultant

Following the Senate Committee on the Judiciary hearing entitled
“Data Security at Risk: Testimony from a Twitter Whistleblower”

on

Tuesday September 13, 2022

1. Senator Hawley asked whether you were aware of communications regarding content moderation between Twitter staff and the U.S. government. You replied, “I’m familiar of the conversations that happened through Department of Homeland Security, the traffic light protocol, where there are messages sent out to organizations about threats that, maybe, the FBI or other organizations had insight to.”

a. Please explain the traffic light protocol, including its purpose, the frequency of use, who sends and receives the information, and how Twitter uses the information.

ANSWER:

From what I saw while at Twitter there were several communications paths with the government but none of the conversations, at least that I was aware of, focused on content moderation.

The methods of communications I was aware of were:

- Legal demands, court orders, subpoenas, etc. Handled by a system called ‘Zipbird’
- Communications from the US Government to US companies containing early warning of upcoming risk and security events.
- Ad-hoc communications with Twitter employees through personal/professional relationships

Zipbird was the official system and process to handle in-bound communications involving legal demands, subpoenas, court orders, etc. The procedures for making these requests are described on Twitter’s website[1]. I am unaware of what percentage of requests were legal demands regarding content moderation. The team managing the initial handling and vectoring of inbound requests lived in Counsel’s organization. There had been discussions to transition the operations aspect of Zipbird to my organization but this had not happened prior to my unexpected departure from the company.

Early warnings of near term threats and risks came into Twitter from CISA (Cybersecurity and Infrastructure Security Agency) and/or DHS. These were emails from USG agencies sent directly to Twitter employees who signed up to receive the alerts. The alerts used the Traffic Light Protocol (TLP) to indicate how the information should be handled and whether it was allowed to be shared within the company or external to the company. Twitter did not, to my knowledge, maintain awareness of who in the company was receiving these government alerts or whether these employees were corresponding with the USG bi-directionally. The Traffic Light Protocol is described in a footnote here[2]. At a very high level RED means no sharing, AMBER means sharing is allowed only within the company, and GREEN would mean the information in the advisory can be shared within relevant industry businesses. There was no request for content moderation to the best of my recollection.

It was also apparent that there were ad-hoc relationships with government personnel in various agencies. These were believed to be personal or professional relationships between individuals. I had laid out plans to better understand, capture, and formalize these communications to ensure they were appropriate. However this effort was immediately scrapped upon my departure. Mr. Agrawal, with 10 years of engineering work at Twitter and elevated to the position of CEO in November 2021, wanted to kill this effort even though the executive team had approved it, the approved effort had been presented to the board, and the effort had been scoped and funded. Given the range and number of suspected relationships within Twitter there was concern about whether these relationships were appropriate and what information was being shared and in which direction (i.e. inbound to Twitter or outbound to the USG). I am unaware of whether these conversations included topics regarding content moderation. However I was made aware that members of Site Integrity, the organization in Counsel with overlapping responsibility for content moderation, maintained some of these ad-hoc relationships.

[1] https://help.twitter.com/en/rules-and-policies/twitter-legal-faqs

[2] Cybersecurity and Infrastructure Security Agency, Traffic Light Protocol (TLP) Definitions and Usage, *available at*: https://www.cisa.gov/tlp

2. You told Senator Kennedy that you have seen numerous examples on underground forums where individuals have offered to sell access to accounts, delete accounts, or unban accounts. Did Twitter investigate the incidents to determine whether the individuals offering these services were employees of the company? If so, what were the results? If not, should Twitter investigate as a matter of course?

ANSWER:

There are several types of Twitter access and services one can find offered for sale in underground marketplaces (sometimes referred to as the "darkweb"). There are advertisements on the darkweb offering to buy "fake followers" (e.g. paying to have a lot of bots instructed to follow your account to artificially inflate the perception of an account being more influential than it is). There are also advertisements offering to sell fake Twitter accounts. These accounts may then be used for a number of purposes (e.g. spam, scams, harassment, etc.). Such fake accounts are advertised as having already fooled Twitter's simple checks to determine whether or not they are "bots".

While both of these are of concern, the most disturbing access being sold on underground marketplaces I was made aware of was "initial access". This meant accounts or credentials that would provide the purchaser with access into Twitter's company and computers.

I believe the Senator is asking about this last class: "initial access".

To the best of my recollection Twitter was not investigating underground offerings of initial access to determine whether the offerings were real or fraudulent.

It was my experience that Twitter would prioritize resourcing and investigating issues primarily only after they had become a publicly known issue and there was negative external press.

Fixing access control, system software and configuration hygiene, and compliance (e.g. data awareness and control), would fix much of this and also make it possible to more easily identify the root of the problem instead of just the symptoms.

Twitter was not mature enough with their internal engineering infrastructure and security monitoring, nor were they staffed appropriately, to try to investigate every underground market offering for initial access (many of which would be scams, though some would likely be valid).

3. Regarding click-through ads, Twitter appears to be aware of the increased risk to user data as compared to non-click through ads.

a. Does Twitter make any effort to protect users or warn users of the potential harm of click-through ads?

ANSWER:

I was not aware of Twitter efforts to educate or protect users from potential increased risk that may be posed through click-through ads.

b. Does Twitter set rules for advertisers on the types of information that advertisers are allowed to collect? If so, what mechanism has been put into place to enforce the rules? If none, please explain why not.

ANSWER:

Twitter's stance was that Advertisers are responsible for their ads, not Twitter. Twitter has a written policy[3] capturing this stance and "asking" advertisers to largely act responsibly and self-police their ads and ad-content. To the best of my knowledge Twitter did not have restrictions or protections in place to enforce safety on click-through ads.

While there are technical challenges to this problem, what I saw within Twitter was that very few people knew what information was being collected or sent to third parties in the first place. Lacking this knowledge more broadly meant it was very challenging for the company to assess threat model risks in this area even if someone had the foresight to consider these risks and concerns about the users of the platform.

c. Does Twitter determine what data a potential advertiser may take from a user before allowing the advertiser on the platform? If not, should Twitter do so?

ANSWER:

Not to the best of my awareness. I heard no mention of this in executive meetings or through my interactions with sales engineers and managers during my time at Twitter. If this was a defined concern it was not a focus that was shared with the executive team. The closest I can recall is Twitter's Ad policies document telling the advertiser what types of content is inappropriate for the service (e.g. weapons, illicit drugs, political content). This document, and largely self-policing policy, does not touch upon technical information that advertisers should limit themselves to collecting from users by way of click-through ads.

4. Your testimony made clear that Twitter is focused on user growth, revenue, and crisis response over data safety. In your opinion, what enforcement actions, whether through Congress or federal agencies, could best motivate Twitter to shift its focus to user protection and data security?

ANSWER:

In my experience one-time fines from the FTC, or any other regulatory bodies, did not meaningfully shift Twitter's focus to user protection and data security. A dedicated privacy engineering team, comprised of privacy experts, was only created after I joined. One-time fines are not viewed as impacting future revenue and hence the executive team, the board, and investors, could treat individual fines as exceptions that would not impact future projections.

[3] https://business.twitter.com/en/help/ads-policies.html

Rules and regulations that represented continuing impact to growth, revenue, or operational tempo, were much more successful at shifting Twitter's focus. Examples include:

- ongoing fines equal to a percentage of revenue
- blocking access to a market (e.g. France, California, etc.) until the problem is corrected
- restrictions on monetizing particular data collected from users (e.g. e-mail addresses, phone numbers, etc.) until the problem is corrected
- mandating new requirements or specifications regarding how internal engineering or business functions and processes must be run going forward
- making members, or a subset of members, of the executive team and upper management have some amount of personal liability (civil and/or criminal)
- Being required to pass audits performed by external companies not paid or contracted by Twitter
- Press about Twitter describing concerns and potential ongoing fines and restrictions that are on the table for continued failure to address problems

5. To the best of your knowledge, are there other companies, besides Twitter, that also suffer from poor data privacy controls or a susceptibility to foreign influence? If so, which ones?

ANSWER:

No. Not to this degree, to the best of my knowledge.

6. During your testimony, you cited the need for increased whistleblower protections for individuals working in the tech industry. What type of whistleblower protections do you believe are needed the most?

ANSWER:

I believe that whistleblower protections should be broad and comprehensive to protect whistleblowers across the private sector. Instead of piecemeal laws that provide coverage by industry or type of wrongdoing, we need a law that will encourage all potential whistleblowers to share information, whether they come from the public or private sector, whether they are employees or contractors, or whether their companies are privately held or publicly traded. Part and parcel to this is to ensure that those protections have teeth because without it, those teetering on whether to come forward or not may elect to stay silent.

Whistleblowers who raise potential securities violations about publicly traded companies have broad protections under the Sarbanes Oxley Act and, depending on the jurisdiction, whistleblowers in the private sector who report unlawful conduct may enjoy strong protections. The federal government should consider similarly broad protections for whistleblowers in the private sector. The FTC Whistleblower Act of 2021 proposed important whistleblower protections for whistleblowers who provide information about potential violations of laws, rules or regulations enforced by the FTC.[4] While this law only covers reports regarding legal violations under the FTC's purview, it includes significant legal components that have proven effective in protecting whistleblowers, so I will highlight a few of those components here that I believe are especially important to include in any general federal whistleblower protection law.

[4] https://www.congress.gov/bill/117th-congress/house-bill/6093

First, the legislation includes a prohibition on retaliation by the employer against whistleblowers who make internal disclosures or disclosures to a government entity. While all whistleblowers are vulnerable to retaliation, tech workers are especially vulnerable because of the wealth and power in the tech industry.

Second, the legislation protects the identity of whistleblowers by ensuring that any identifying information about the whistleblower is not subject to public disclosure. At the same time, for whistleblowers who wish to come forward publicly, it allows them to do so by prohibiting the enforcement of arbitration agreements for these specific claims.

Finally, the proposed legislation also includes a component that is important to encourage whistleblowers to come forward, despite the tremendous personal, financial, and career-related risks—an award program. Award programs for whistleblowers ensure that whistleblowers who face harm to their careers and reputations will nonetheless be compensated for providing information that is relevant and significant for any enforcement action.

The tech industry is only growing, so ensuring that workers in this industry are protected for reporting various types of unlawful conduct is essential to protecting consumers, users of these services, and the public.

7. Based on your testimony related to Twitter's management of users' personal data, in your opinion is it possible for Twitter to fully comply with (1) the California Consumer Privacy Act and (2) the General Data Privacy Regulation's individual right to request that someone's personal data be deleted as required by these laws?

ANSWER:

At the time of my employment it was not possible for Twitter to be compliant with a request that their user data be deleted.

The company had known for over 10 years that they did not know where user data lived within their systems and who had access to it or how it was protected.

Because of this, I did not understand how the company could be compliant for subpoenas that demanded all data that Twitter had on a particular user.

I found many talented engineers and passionate employees wanting to fix these problems at Twitter, executive leadership at the company lacked cohesion, experience, and the expertise required to close on the underlying issues.

The Privacy Engineering organization was created after I joined. I helped bring on one of the most capable Privacy Engineering Leaders in the US. It was only at this time that the first formal measurements of the scope of the data problem were taken. It was only after such measurements could the depth and extent of the problem be known and hence a plan and end state defined. Without having measured the problem in the past, I do not believe it was possible for the company to have been in compliance with CCPA or GDPR.

Towards the end of 2021 I began including these numbers and the severity of the problem in almost every executive team meeting. Several executives commented that they were aware that this problem existed and they were not surprised that the company had not made appropriate progress for years, yet still were unwilling to devote necessary resources in their organizations to bring data their teams were producing into compliance.

8. What would it take for Twitter to fix its current inability to know the full universe of data, personal or otherwise, that it maintains and to determine where it is stored?

ANSWER:

Every company accrues technical debt as it gets going. As they become more established, it is important that they periodically pause to go back and pay those bills, or else that debt becomes compounded. At the executive level Twitter has intentionally ignored the need to perform this kind of house cleaning for years, which has led them to their current state of disorder.

With the existing service it would be necessary to have executive support, priority, and resourcing across the board. Leadership, both executive and senior managers, would need to have prior experience and operational success in performing these turn-arounds. All leadership would need to be data driven with a significant effort put into active and ongoing visibility into the systems and processes related to these issues.

I described the areas of critical concern that would need to be addressed in the document I wrote and sent to the Twitter Risk Subcommittee of the Board of Directors. I also shared this road map with the new CISO at Twitter, whom I spoke with shortly after my unexpected termination.

The problems at Twitter were not new problems to the industry. The problems were the lack of basics and fundamentals. Basics and fundamentals that most companies engineer away at the beginning or that they revisit and fix early on in their lifespan. Twitter did not do this and built on top of significant deficiencies for over a decade. Even with the talented engineers at Twitter, they had been unable to address the root of the problems. This is because of aversion to impacting short term returns by the executive team, lack of expertise and experience at the executive, board, and senior leadership levels, and a culture that encouraged glossing over problems and being driven by crises.

While many of the problems are quite basic in nature even with the correct appetite, executive support, experience, and culture, this is a multi-year effort to sufficiently address. The company is in the position of needing to make up for 10 years of neglect and debt accrual and that can be compressed only to a particular extent while still maintaining and running the service.

Another option is to rebuild the Twitter service from the ground up and then switch over to the new service. This may sound extreme, however senior engineers at the company performed evaluations and estimated that it could be easier, faster, and less expensive to rebuild from scratch, addressing these issues in the process, and switching over once completed. Other companies have taken this approach in the past.

Irrespective of which path is taken Twitter must get a grip on what data they have, why they have it, under what context it was created or collected, how the data needs to be protected, when it needs to be deleted, what systems and people operate upon or touch the data, etc.[5][6].

[5] Such systems include custom built systems, databases, and filesystems (e.g, HDFS).

[6] The fact that Twitter was unaware of where data was being accessed and used after it was collected was demonstrated when this deficiency was identified by the French CNIL. Twitter had to address

In order for Twitter to become compliant with the GDPR or other data use requirements new data coming in needs to be handled differently and the existing data needs to be mapped, identified, and brought into compliance with regards to privacy and security. This means identifying, understanding and changing Petabytes[7] of data Twitter has and the modifying both the data and the systems that handle the data[8].

9. If a Twitter employee had access to main systems and inappropriately accessed user data, would Twitter have any way to know that this occurred, what data was accessed, or what was ultimately done with that data?

ANSWER:

In general no. There was insufficient logging (and/or insufficient monitoring of logs), a lack of awareness of data, and inappropriate access control. While there may be certain situations where Twitter could know these things they were the exception rather than the norm. I feel confident in this response because of multiple times where it was necessary to understand what had happened on certain systems, or who had accessed or created particular data and I was repeatedly informed that it was unknown and that there were no ways to figure out the answer to such questions.

10. What would it take for Twitter to address the employee access vulnerabilities to better protect personal data and data that employees don't need access to in order to perform their job duties?

ANSWER:

As stated elsewhere in these responses I was made aware by senior engineers at the company that to address these (and other critical issues) it would be faster, easier, and cheaper to rebuild the Twitter service from the ground up and then switch over to the new service than it would be to retroactively address the decade of technical debt and design choices in the current system.

For Twitter, as of when I was terminated, one of the key problems in need of solutions was knowing what data they have, where it lives, how it is processed, and how it needs to be processed and accessed, and by whom. Without this, it is impossible to attain and confirm compliance, security, and privacy requirements and goals.

non-compliance with privacy involving "cookies". Cookies are pieces of data that Twitter provides to users' web browsers to keep track of the users and activities. Executives and engineering alike at Twitter did not know how they used their own cookies and what systems depended on them. This alone meant that Twitter was likely not in compliance with GDPR. Because of Twitter's lack of awareness and understanding about how their own service ran, attempting to trace cookies through the Twitter systems turned out to be impossible. Twitter had to ultimately resort to changing cookies in their live system and then waiting to see which of their backend systems subsequently broke because of the change.

[7] As a sense of scale a single Petabyte is equivalent to 20 million 4-door filing cabinets full of documents and Twitter has 100s of Petabytes of data.

[8] In 2018 it was reported that Twitter put 300 Petabytes of data in Google Cloud for off-line analytics. As of January 2022 Google Cloud was only one part of Twitter, not responsible for the actual running service which was still being run only in Twitter's data centers. (https://www.lightreading.com/enterprise-cloud/data-strategy-and-analytics/twitter-moves-300-petabytes-to-google-cloud--thats-a-lot-of-covfefe/d/d-id/746167)

11. You said that Twitter was essentially allowed to self-grade their compliance with the 2011 FTC consent decree. Do you have suggestions about what the FTC should have done to ensure Twitter complied with the consent decree?

ANSWER:

The FTC should require technical elements and evidentiary data to back up claims and responses made in company reports. With requests and questions less ambiguously defined, the FTC would be in a better situation to know that answers represented the actual state of affairs and were not cleverly worded responses[9]. This would have meant Twitter would need to actually have made broad changes as intended/required by the FTC, or else lie outright.

There was a lack of data driven scrutiny by the FTC, or perhaps insufficient understanding, to stress-test answers provided by Twitter in response to inquiries. Based on my experience at Twitter, the appearance of compliance was often achieved by word play or by crafting non-answer answers to questions posed that was possible because ground truth data was not demanded. For example, I was made aware that when the FTC asked questions regarding a particular technical deficiency, Twitter would present a hyper-specific example of mitigation. What was presented as a broadly applicable response was in fact an exception and not the norm. The wording and non-representative examples would create a misrepresentation, without technically saying an outright lie in how the answer was stated. Requiring data that can be verified to be truthful and that can be verified as actually representative of state of affairs across the whole company would address this.

The FTC could also employ outside technical auditors. These auditors should be independent, without relationships to the company the FTC is dealing with. This way the information being returned would not come from, or be influenced by, the company being evaluated or investigated. There are perverse incentive structures at play when the company being investigated is paying the company charged with performing the evaluation.

12. Congress is currently considering federal data privacy legislation. Do you have an opinion about the American Data Privacy and Protection Act?

ANSWER:

I am not a lawyer or a legislator so I don't consider myself qualified to opine on legislative proposals. However, I am in favor of a federal privacy law and believe certain fundamental principles should be part of any proposal that Congress might consider. I believe federal privacy proposals should:

* Require robust protections for personal data, including limitations on the purposes for which personal data can be collected, used, and transferred without a person's affirmative consent.
* Require privacy by design and implementation of up-to-date security practices appropriate to the sensitivity of the particular data collected and stored.
* Give people rights to access, correct, delete, and port their data elsewhere as they wish.
* Ensure appropriate agencies have the authority, resources, and enforcement mechanisms necessary to take action against those who violate the law.
* Make it necessary for companies to have audit logs for data that spans the life of the data.

[9] Such statements potentially accompanied with isolated examples that may or may not represent the larger situation truthfully.

* Have up to date statistics and data that can be requested and required at any given time that shows the correctness and totality of data privacy and security.
* Demonstrate awareness of private citizens' range of interests and concerns so they can protect their own interests appropriately.
*Not preempt (and thus eviscerate) stronger state privacy laws that may be in effect.
* Avoid creating situations that would prevent states from taking steps as they see fit to protect privacy in the future.
* Be flexible enough to adapt to future innovation, or else privacy protections will quickly become outdated and ineffective.
* Require answers and attestations that are driven by data and that can be independently verified (not allowing companies to self-certify, or otherwise "grade their own homework")

13. Some argue that federal data privacy legislation is not necessary and that companies can self-regulate. Do you believe that companies would effectively self-regulate based on your experience?

ANSWER:

Unfortunately, no. Based on my 30+ years of experience and my knowledge of technology companies, the incentive structure leads to a deprioritization of privacy, security and public health and safety that does not strike the appropriate balance between profit and security/privacy.

14. You allege in your disclosure that Twitter is knowingly infringing on intellectual property owned by others. How long has Twitter been intentionally infringing on intellectual property owned by others?

ANSWER:

I was informed that this situation had been previously raised to the executive team and to the Board by the Chief Privacy Officer ("CPO") in years past, that the issue was acknowledged and understood, but that no action had been taken. This was brought to my attention only a few days before my abrupt termination and thus I was unable to investigate further.

a. To the best of your knowledge, were the owners of the intellectual property currently being infringed by Twitter aware of that infringement?

ANSWER:

As this was brought to my attention only a few days before I was abruptly terminated I was not able to find out these details.

b. To the best of your knowledge has Twitter ever intentionally engaged in discussions to license intellectual property and then subsequently infringed that intellectual property instead?

ANSWER:

As this was brought to my attention only a few days before I was abruptly terminated I was not able to find out these details.

QFRS
Senate Judiciary

Questions from Senator Tillis for Peiter “Mudge” Zatko

1. **According to a September 9, 2022 Wall Street Journal article Twitter shareholders are being asked to vote on Elon Musk’s proposed $44 billion takeover of the social-media company on the same day as this hearing. What are your thoughts on this proposed purchase and the timing of the vote?**

ANSWER:

I take no position on the proposed purchase or the timing of events associated with the proposed purchase.

2. **It has been reported that Twitter might have infringed on intellectual property rights regarding internal use of various software tools. What details can you provide regarding this matter – in your opinion, how could this sort of oversight occur in such an established tech-based company such as Twitter?**

ANSWER:

I was told that Twitter did not have the appropriate licenses for the training data they used to create core machine learning models. It was my understanding that these models were a key component of the service and that if Twitter were instructed to stop using these models it would be detrimental to the service and company. I was told that this situation had been raised to the executive team and to the Board in prior years and that both teams acknowledged the issue but that no action had been taken.

As all of this was brought to my attention only a few days before my abrupt termination I was unable to investigate further.

As to my opinion on how this sort of issue could occur in a company such as Twitter, based on my experience Twitter is a company that is in a constant state of reacting to one crisis after another. Because of this issues would get dropped before they were appropriately completed in order to handle the next crisis. What I saw while working at Twitter were many new crises that were actually the result of not correcting and completing fixes of a previous crisis. This pattern of behavior does not lead to meaningful long-term solutions and could be relevant to answering your question.

3. **There have been several instances reported of different foreign governments and agents attempting to influence Twitter and gaining access to sensitive user information.**

   a. **What steps did Twitter take, whether successful or not, to respond to these attempts when you were employed there? And in your opinion where does Twitter stand on this topic today?**

   ANSWER:

In my experience Twitter was largely reactive and as such would discover incidents of foreign influence and infiltration either by having them externally identified and reported, or by discovering internal issues by accident. It may have been that mine was the first case where proactive discovery and identification of foreign agents inside Twitter was achieved proactively and with intent.

b. Specifically, what kind of information could these governments and agents have obtained? And in your opinion what sort of dangers would this pose to governments and individuals?

ANSWER:

There are numerous ways for an agent inside Twitter to provide value to their external "handlers" and there are numerous types of data that may be of interest. For instance non-public information about users could be used to reveal real identities of dissidents for the purposes of harassment, intimidation, persecution, or execution. Certain data could be used to confirm real-time geolocation and activities of a targeted user for coordinated external activities. An agent could obtain intelligence about technical limitations of the platform to provide intelligence and guidance for other espionage activities that are intended or ongoing on the platform.

Equally valuable to a foreign government is information about what tactics and overt pressure that a country is placing on Twitter is having influence on Twitter internally. For instance this would be valuable in understanding whether threatening harm to employees in the country was having meaningful impact in Twitter leadership decision-making around censorship or business activities in various parts of the world.

Countries running disinformation operations on the platform may want to know whether Twitter had internally identified these operations, or similarly whether other country's spies and activities were well known internally and if actions were intended.

An agent at almost any level and role in the company would have been able to report back to their foreign handlers that Twitter was largely incapable of identifying compromises and activities of state actors that would, or already had, compromised the company's infrastructure. This would provide a green-light for new cyber compromises or to inform intelligence communities that they could continue operations largely without fear of discovery.

It had already been demonstrated that it was possible to co-opt accounts of powerful people once internal access to the company had been acquired.

4. The independent Alethea Group report, which you asked for while employed at Twitter, disclosed that Twitter's team responsible for enforcing site integrity policies were understaffed and relied on external reports for its counter-disinformation effort.

a. In your opinion, to what extent was Twitter able to timely, accurately, and adequately respond to disinformation in light of these reported issues within the company?

ANSWER:

Twitter's efforts and ability to respond adequately was minimal, superficial, and sporadic rather than systematic. The abilities were primarily English language centric and largely done reactively rather than proactively.

Abilities to ensure integrity during US elections, as captured in the report, were primarily achieved by manually staffing people to review. This meant that the solutions for "election squads" were not scalable, significantly non-automated, and non-transferable to other parts of the world due language constraints of the humans performing real time reviews.

b. To your knowledge or in your opinion, to what extent did Twitter rely on the Biden administration in deciding what news stories were considered disinformation?

ANSWER:

I was unaware of this occurring at Twitter while I was there.

5.

a. Based on the reported lax or nonexistent security protocols within Twitter, what negative impact did you see or could you foresee on elections such as voting?

ANSWER:

I would expect a continuation of challenges at Twitter that the public has already seen and that were touched upon in the independent thirdparty analysis document I attached with my disclosure.

b. And more broadly, what sort of censorship did you see or do you foresee as being possible?

ANSWER:

I saw a company that was almost always running flat out attempting to react to the latest crisis that had popped up. I personally did not see intentional censorship at the company.

As for unintended and unwanted censorship, there was constant pressure from foreign governments. This pressure, especially when involving potential hostility or safety risks to Twitter employees, was having success in driving the head of Counsel to consider complying with such demands.

Discussions were had about the possibility of ceding censorship capabilities to other countries, some of which are objectively known to be un-democratic.

The lack of visibility into systems, processes, and data, makes it difficult to make an attestation of perfect correctness of operations.

Finally, while I was there I did not observe priority focus across executive team meetings in regards to tracking moderation bias that could capture and identify patterns of censorship.

c. **To what extent could a Twitter employee create new content under a user account that they do not own or alter posted content under an account that they do not own?**

ANSWER: [capability of posting from accounts]

I was informed by numerous engineers that it was possible for anyone with basic engineering access to Twitter's production environment[10] to figure out how to find the right data and systems in production to allow them to tweet as anyone on the system. in production where they could directly access data and post content as any user on the platform.

6.

a. **Based on the report that Twitter engineers worked off of live production data and tested directly on the commercial service, as opposed to first using a test environment, in your opinion how does this compare to the practices of other established tech-based companies?**

b.

ANSWER:

This is a significant indicator that Twitter has fundamental technical deficiencies that are normally addressed early on in a company's lifetime. Twitter is an outlier in this deficiency, as well as others, in comparison to peers.

c. **What in your opinion could be the possible reason or reasons for not utilizing a test environment?**

ANSWER;

From what I saw, the following could all be contributors:

Not taking the time and resources to pay off technical debt and instead allowing it to continue to accrue year after year.

Lacking executives with sufficient knowledge and experience across companies in the industry to understand common industry practices.

Having a culture and environment that does not measure and reward strong execution in basic Run the Business / keep-the-lights-on (RTB/KTLO) efforts; having a culture that is excessively confrontation averse and is reluctant to highlight and address technical deficiencies; having a culture that is incentivized to hide problems, celebrate isolated wins, and to not engage in constructive critiquing

Lack of accountability at all levels

Not being fundamentally driven by data

Lack of dashboards with relevant data and context visible at the executive level

[10] This access was provided to every engineer when they joined the company.

Being a company that is driven by crisis or that needs a crisis in order to execute

Teams and groups operating in silos and resistant to cross team collaboration

d. **Where there any attempts while you were at Twitter to change this practice of not using a test environment?**

ANSWER:

Yes but to little effect. The two people most vocal in championing the creation of a test environment to remove the risk and need for engineers to have access to production were myself and a VP of revenue engineering.
I identified this as a priority. However, as addressing this issue would require significant effort and long overdue changes to address technical debt, the effort repeatedly met with both technical and non-technical resistance. There were some small efforts to approach this problem but it was not viewed as a priority in the Engineering organization.