

Issues and Objections Regarding Twitter InfoSec Information and the Q4 2021 Twitter Risk Committee

Peiter Mudge Zatzko

Latest Revision: February 2022

Overview	2
Introduction	2
Part 1: The Problem	2
Incorrect and Misleading InfoSec Information: Q4 Risk Committee	2
Identification of the issues in 2021	3
What is Twitter's accurate InfoSec situation, as of Q4 2021?	6
This did not happen overnight	7
Part 2: A More Truthful Q4 Description of Twitter Key Information Security Risks	8
Overview	8
4 Areas of Critical Focus	9
Access Control	9
Access Control and Insider Risk	11
Security Patches and Software Configuration and Versions	12
Client fleet (laptops)	12
Servers (data center)	13
Processes and Compliance	15
Incidents	16
Part 3: Inaccuracies in InfoSec Materials Presented Q4 2021	18
The Deck	18
Access Control	18
SDLC, Security Reviews, Privacy Reviews	21
Patches and Software Configuration and Versions Compliance	23
A note on Zero Trust and Endpoints (Employee computers)	25
Incidents and Incident Classes	25

Overview

Introduction

This document describes (Part 1) events leading up to the transmission of inappropriate information to the risk Committee over objections of the Twitter Security Lead. It then provides a (Part 2) Replacement InfoSec Risk Report on Top Risks. The document closes with (Part 3) select descriptions of inaccuracies and misrepresentations in the materials that the CEO instructed be presented to the Q4 Twitter Risk Committee. Electronic records support this description of events and the information contained herein.

Part 1: The Problem

Incorrect and Misleading InfoSec Information: Q4 Risk Committee

In December 2021 the Risk Committee received information about Twitter's information security posture that is inaccurate and misleading. It appears that Twitter's information security environment has not been accurately characterized to the Board of Directors and Risk Committees dating back to before my tenure. This disconnect may exist elsewhere¹ but the focus of this document is on Information Security.

It is critical that the Board have accurate and truthful views into Twitter's InfoSec issues and posture. The Board needs this accurate information so they can take corrective actions and ensure reports to regulatory and other bodies are accurate. It is crucial that the Board is not misinformed as Twitter works to comply with the existing FTC consent decree, faces possible new violations related to mis-representations in ongoing consent

¹ There was also a disconnect between Twitter's stated Privacy posture and the reality of Twitter's privacy issues. However, by removing Privacy Engineering from Information Security and through the work of myself, [REDACTED], and the new team, this disconnect has been significantly improved. Several ICs have volunteered to provide questions that members of the Board should ask to determine where else this disconnect and misrepresentation may, or may not, exist in Twitter.

negotiations, new demands from other regulators², and moves into more regulated environments³. If the Board is misinformed, representations and statements to the outside world will be inaccurate as well. This would impact Twitter users (customers) and shareholders.

One of the reasons I was hired was to evaluate Twitter's information security environment and provide an accurate assessment. There were concerns that Twitter had serious problems in these areas, which threatened its data security and integrity in its industry. With my domain expertise I was to look into these concerns and make the accurate and truthful state of Twitter's security posture known to the executive team, and the Board, and work to put Twitter on the right path forward.

Identification of the issues in 2021

██████████, ██████████, ██████████, ██████████, ██████████. I raised concerns to Human Resources Business Partners (HRBPs) and I spoke to ██████████ numerous times about my concerns regarding accurate assessments of our security and the necessity to report accurately to the board and its subcommittees on security gaps and where improvement was needed. ██████████ walled me off from Twitter Information Security data and those who were in ██████████ organization. ██████████ actively resisted my direction. ██████████ prepared a report to the Q4 Risk Committee, which report was a whitewash of security problems at Twitter and overstated the improvements made to security. This report contained inaccurate and false statements which I believe were intended to mislead the Board. My efforts to rein in ██████████, whom I viewed as neither sufficiently competent at ██████████ position nor forthright, were opposed by Human Resources due to Twitter's fear of a lawsuit based upon ██████████ many allegations, which were all investigated by Employee Relations and found to be unsubstantiated. I believe these were made to deflect criticisms of ██████████ work.

I concluded in October 2021 that ██████████ should be offboarded due to performance issues. Human Resources informed me that the documentation and communications between myself and ██████████ and the documentation of ██████████ failure to perform met the requirements to offboard ██████████ and that it would be done. At that point Employee Relations was permitted to share with me that they had received numerous other complaints about ██████████ inappropriate behavior.

² Irish DPC, French CNIL

³ E.g. Money Transmission Licenses (MTL)

██████████ offboarding was then repeatedly blocked, including directly by Parag Agrawal. Had ██████ offboarding occurred when approved, going back to October 2021, ██████ report would not have been released to the Q4 Risk Committee.

When I brought to Mr. Agrawal's attention the fact that ██████████ report was misleading, inaccurate and intentionally wrong, he overruled me and overruled my recommendation that ██████ report be rewritten to make it accurate. Mr. Agrawal ordered me to permit ██████████ report to the Q4 Risk Committee's December 16, 2021 meeting with the instruction that I walk back the many and fundamental inaccuracies and falsehoods contained in ██████████ report after ██████ presented it to the Committee. After registering my concerns about this approach on the record, I followed Mr. Agrawal's instructions as best I could. Mr. Agrawal committed to assist after-the-fact in correcting the record. However, after the Q4 Risk Committee meeting, Mr. Agrawal expressed disappointment that I had not completely walked back the report. What he refused to recognize was that walking back a report that instead needed to be repudiated was an impossible task. In essence, he chided me for not telling the Q4 Risk Committee to completely disregard the report submitted to it.

In addition to e-mails that I sent to Mr. Agrawal, and to the head of HR, prior to the Risk Committee meeting, and those immediately following the meeting, I continued to communicate that the situation was not resolved. In receipt of one of my emails, January 4, 2022, where I repeated concerns about the false representations that were made to the Q4 Risk Committee at its December 16, 2021 meeting, Mr. Agrawal replied (January 6, 2022) that he was "surprised" by the allegations I made. These allegations were merely a recapitulation of my prior complaints and concerns.

My reporting of false information triggered an Audit investigation regarding misrepresentations and false statements being made to this subcommittee of the Board. Approximately two weeks later, on January 19, 2022, my employment was terminated. One day prior to my termination Mr. Agrawal held a surprise meeting that included the Head of the Risk Committee. At that meeting Mr. Agrawal falsely stated that he had ordered me to redo ██████████ report a month prior and was still waiting for the corrected information. The electronic record will verify my recollections over his.

The termination of my employment appears to be in direct violation of New Jersey's whistleblower protections (Conscientious Employee Protection Act - CEPA, N.J.S.A 34:19-1, et. seq.). Regardless of my unlawful termination, I *still* view it as crucial that Twitter's accurate and truthful measurements of Twitter's security risks and posture be correctly conveyed to the Risk Committee and the Board.

Below is the accurate assessment of the state of Twitter's current security, as should have been conveyed to the Q4 Risk Committee of the Board.

– Twitter is grossly negligent in several areas of information security.

– If these problems are not corrected, regulators, media, and users of the platform will be shocked when they inevitably learn about Twitter's severe lack of security basics. They will lose confidence in Twitter and this will have real world impact to the platform and to the company.

– Regulators, when evaluating Twitter, *will* identify these as **systemic issues**. They will likely levy new fines and/or increase existing fines. They will also impose constraints and requirements on how Twitter operates, constraining Twitter's freedom to choose how it executes in various areas of engineering and what Twitter chooses as priorities. Further, Twitter may be precluded from conducting business in certain markets.

There are 4 critical areas that have not been accurately represented to the Board:

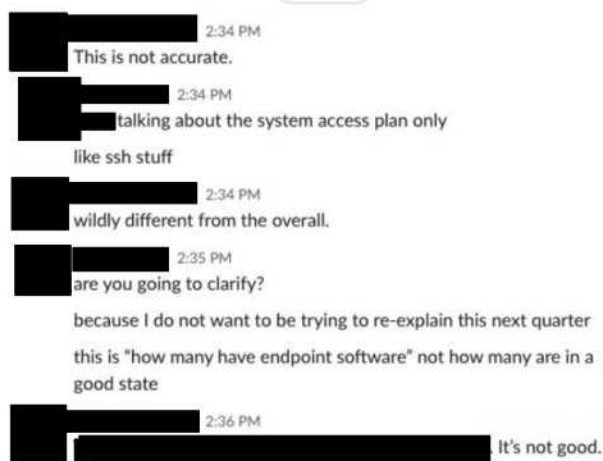
- Out-of-date software and the lack of basic security configuration in existing software (Software and Security Versions/Configurations/Patches)
- Gross problems around access control to systems and data (Access Control)
- Lack of basic processes and compliance such as software development lifecycles, line-managers being allowed to unilaterally overrule security and privacy findings, and a prioritization of running products with known violations over compliance with regulatory requirements⁴ (Processes and Compliance)
- A volume and frequency of security incidents impacting a large number of users' data that is frankly stunning (Incidents)

Before I provide an accurate description of Twitter's Information Security and a critique of the data that was provided to the Risk Committee, it is worth pointing out that other senior people also started to identify that inappropriate information was being presented.

⁴ A recent example includes Twitter choosing not to comply with regulatory requirements, even though it could, until it could optimize more profit within a single region. (French CNIL)

Listening to [REDACTED] abbreviated verbal presentation to the Q4 Risk Committee, both the Twitter Chief Privacy Officer and Twitter's Distinguished Privacy Engineer (the highest engineer rank, equivalent to a VP) objected to what they heard.

[REDACTED] made statements about access control at Twitter and then about endpoint (employee computers) security health. The following was an unsolicited live response to what was said:



2:34 PM
This is not accurate.

2:34 PM
talking about the system access plan only
like ssh stuff

2:34 PM
wildly different from the overall.

2:35 PM
are you going to clarify?

because I do not want to be trying to re-explain this next quarter
this is "how many have endpoint software" not how many are in a
good state

2:36 PM
It's not good.

Slack messages from the Chief Privacy Officer and Distinguished Privacy Engineer refuting statements made by [REDACTED] about Access Control and Endpoint security (in the above messages the topic changes from accuracy of Access Control to accuracy of Endpoint security at "this is 'how many...").

What is Twitter's accurate InfoSec situation, as of Q4 2021?

Twitter is **very** far behind the industry in key areas of Access Control, Software and Security Patches/Configuration/Versions, and Processes and Compliance. This is evidenced in the volume and frequency of Incidents. In more than one of these areas Twitter is a decade behind peers such as Google and Facebook.

Some newsworthy highlights are that more than half of Twitter's 500,000 servers are running out-of-date Operating Systems so out of date that many do not support basic privacy and security features and lack vendor support⁵. More than a quarter of the ~10,000 employee computers have software updates disabled! More than half of Twitter employees have access to Twitter's production environment – unheard of in a company the age and importance of Twitter, where nearly all employees have access to systems

⁵ E.g. encryption at rest, kernels, updates, etc.

or data they should not. At Twitter engineers work on live data when building and testing software because Twitter lacks testing and staging environments; work is instead conducted in production and with live data. With this understanding, it is somewhat less surprising that frequent security incidents are so commonplace at Twitter that more than one per week, on average, occurred in 2021 and were determined to involve millions of people's accounts/data. This additionally provides plausible explanations for some of the numerous platform disruptions, as engineering errors that happen during testing occur in, and impact, production.

This did not happen overnight

To get where Twitter is today took more than just a lack of prioritization on areas of information security and privacy across the past year. This took many years. To get to Twitter's current state of insecurity required repeated downplaying of problems, selective reporting, and leadership ignorance around basic security expectations and practices.

Part 2: A More Truthful Q4 Description of Twitter Key Information Security Risks

Overview

Twitter is nearly a decade behind the industry (and peers) in access control. It is significantly behind publicly traded companies in keeping servers and clients up to date with software and patches. Twitter lacks necessary visibility into networks and systems that it needs to state confidently whether identified security problems have been remediated to the extent necessary, and Twitter experiences an outsized frequency and volume of incidents above that of other companies⁶. The incidents are unsurprisingly rooted in the areas of greatest deficiency.

Internal dashboards show 30% of the ~10,000 employee computers reporting that they are not correctly configured to accept software updates⁷. 60% of servers in the Twitter data centers are running out of date (even unsupported) Operating Systems⁸. More than 50% of full time employees have access to Twitter's production environment⁹ because Twitter does not have appropriate development and testing isolation. And Twitter dealt with more than ~50¹⁰ incidents in the past year stemming primarily from systemic areas of risk such as access control.

These truthful views presented here were obscured by Twitter's bias towards presenting individual "wins" to the Board without larger context. Such misrepresentation creates an impression that improvements have been made when in fact core problems have grown. This is a habit that appears to be longstanding within Twitter and is not dissuaded at the senior, and executive, levels.

What the Risk Committee should look for: presentations to executives, and from executives, containing numbers without context. Such information is difficult to interpret. Numerators without denominators lack important context, and as a result are often misleading. The first question should be "out of how many?", followed by "is this

⁶ This statement comes from my experience at Google, Stripe, Motorola, and InfoSec visibility across dozens and dozens of corporate environments during my 30+ year career.

⁷ Uptyx

⁸ Platform Engineering Dashboards - Kernel Compliance/Non-Compliance and Software Compliance/Non-Compliance

⁹ The source for this statement comes directly from Twitter's LDAP server. This server contains employee access rights to systems and resources.

¹⁰ My notes capture 48 formal security incidents from April through November 2021. Additional incidents during {Jan, Feb, Mar, Dec} 2021 are extrapolated.

enough? Where should we be? What trends were revealed and what was the value? How much did a contextless number, such as a number of security reviews lacking the number of projects needing security reviews, interrupt or delay business efforts and cost resources such as time and headcount to the company and projects?”. We will discuss this problem further in the third section of this document where we discuss specific inappropriate representations in the Information Security presentation.

4 Areas of Critical Focus

There are **4 critical areas** of information security that ***have to be the focus*** of the Twitter Information Security Team and tracked by the Risk Committee: *Access Control, Processes and (Regulatory) Compliance, Patches and Software Configurations and Versions, and Incidents*. Unless something exceptional happens elsewhere, the focus of reporting to the Risk Committee should not be pulled away from these fundamentals until they are addressed. These areas, it turns out, have not been accurately described in the past.

Access Control

Access Control - Twitter is an outlier in this area of risk and not in a good way. Most companies work to restrict access to production systems to only a small handful of people because production systems contain extremely sensitive data and issues in production directly impact customers. Engineers at other companies do their work in testing and staging environments, strongly isolated from the crown jewels of production systems that provide the actual running service. By comparison, Twitter engineers and developers perform work directly in production or interfacing directly with production systems and data. Twitter is where Google was prior to 2005-2007, when they identified and addressed the key issue of removing broad employee access to their production environment. Most companies recognize the risks of outages, sensitive data access, and maintaining the integrity of their platform, and intentionally remove almost all direct interaction with production systems and prevent engineers and employees from having access to production data.

Contrary to what may have been heard or read: *Twitter's access control risk is growing, not shrinking.*

At the beginning of 2021, 46% of all FTEs had privileged access to production systems and data. By Q4 2021 this number was 51% of employees. Twitter has grown meaningfully in its number of employees. The percentage of employees with privileged access has increased on top of this.



Access to Twitter's Datacenter Production Environment¹¹

- i. Dec 2020 46% of employees (2,763 out of 5917)
- ii. Dec 2021 51% of employees (3,995 out of 7714)

(*) The dip was an unintended (internal) incident

Companies following basic security principles do not allow this type of access to production systems and live data. Companies have long ago learned to separate production systems and data from testing and staging environments. Twitter does not do this. Twitter's level of exposure and risk in this area far exceeds the industry.

There are smaller pockets of FTEs who are bestowed even further sensitive access. For instance, they can power on/off computers in the data center or they can perform administrative access on servers in Twitter's data centers. It is important to remove or reduce the access these groups have but they are also edge cases. They are a subset of the primary exposure.

Why should these edge cases not be the sole focus? Consider the understood way to predict the likelihood of an unwanted event from a specific risk, like access control. Assumes a random compromise (or malicious behavior) of a Twitter FTE's account. There is a 1 in 2 chance that the compromise happens to one of the 3,995 out of 7,714

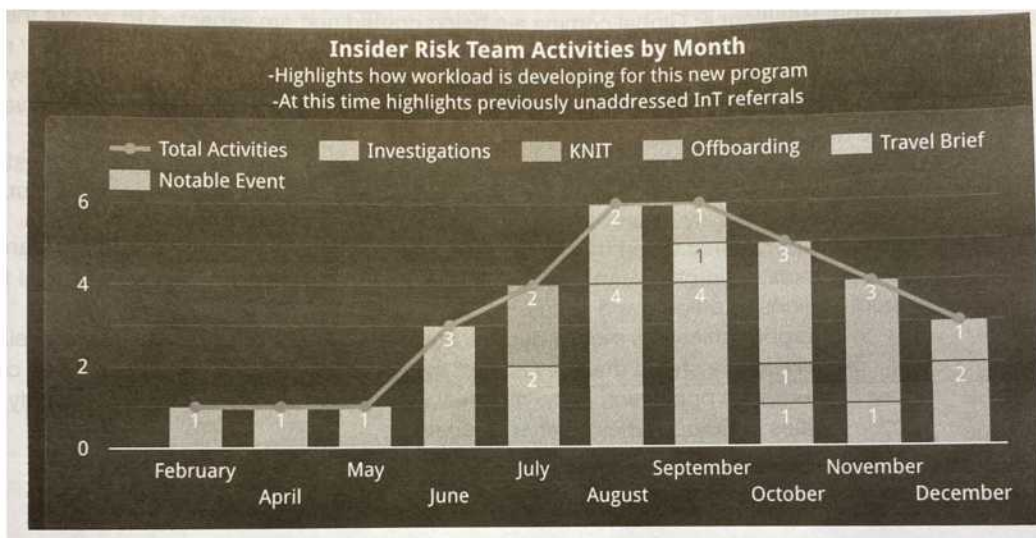
¹¹ It is important to note that this access issue impacts international expansion and operation. Any country where Twitter has an engineer, there is access to production systems and data. Other companies have development, staging, and test environments to mitigate this risk, but Twitter has none of these.

accounts with privileges providing access to production and production data. Now consider the risk for a group of FTE accounts with further access within production. There are ~300 such accounts with god-mode access in Twitter. These represent a 1 in 33 (~300 out of 7,714) chance that the random compromise lands within this edge case.

This is not to say the edge cases are unimportant. But, removing these edge cases without a plan and processes in place to allow work to be done outside of production will prove to only be temporary improvements. This was seen in the repeated focus on reduction of this exact edge case of access control at Twitter. Without a plan for being able to reduce broad access the reduction of this small subset almost immediately began re-growing after the initial reduction. The people whose access were reduced needed it back to do their jobs. This subset reduction is discussed in more detail in section three because this subset reduction was presented, without important context, and stated in a way to imply the larger production problem was being solved when it was not.

Access Control and Insider Risk

There are several known insider threats (KNITs) at Twitter¹². Because of the ubiquitous access to production systems and/or data and the lack of isolation environments and logging, this risk is significant. Combine this with ~30 offboardings per week, each of which represent periods of enhanced concern for insider threat, and the lack of access control and ubiquitous access grants are critical problems.



¹² Referenced Q3 Risk

This chart shows the beginning of Insider Risk Tracking (Offboarding pace via JIRA is about 30/week; most of which are not adequately tracked for insider risk) - this is an improvement from Twitter's lack of Insider Risk abilities in Q1 2021. This chart, and effort, is primarily maintained and run in the Corporate Security Organization.

Twitter has limited ability to effectively constrain and mitigate insider risk without having mature access control and a separation of sensitive data and systems beyond what Twitter presently possesses. Correct access control is also a critical path item for privacy. It is required for regulatory and compliance and to meet expectations and representations made to users and the public. As will be shown below, Twitter systems and servers lack basic security compliance, making this even worse.

Contrary to what ██████ told the Risk Committee in the December meeting, at present there is not an agreed-upon plan to address the broad access control issue. In place of a plan there is a *goals-focused* document within the Twitter Information Security organization but Engineering, Privacy, and IT have not signed off on the approach and there remain significant questions around the feasibility of Information Security's understanding and approach around the effort.

Security Patches and Software Configuration and Versions

Client fleet (laptops)


When ██████ told the Risk Committee that nearly all Twitter endpoints (laptops) have security software installed, this statement lacked the following critical context:

Almost 40% of these ~10,000 employee computers (aka endpoint systems aka the client fleet) are not in compliance with basic security settings.

30% of the total endpoint systems report that they do not have automatic updates enabled.

These are the systems used to access Twitter's source code, internal systems, and sensitive data. It turns out that this unacceptable state of security on employee systems has not improved from Q1 2021. Throughout 2021 ██████ made numerous statements and references to myself, and others, that the client fleet was in good shape. After all, ██████ stated, most laptops had some security software installed. What the security software reported was either not understood by ██████ or it was understood

but it was chosen not to be made a priority focus. The following is what the software revealed.



Specific details and numbers can be seen on Twitter's Uptyx dashboards¹³


██████████ first actions should be identifying, reporting, and addressing this problem. With focus, it is a 3-5 month project to reach acceptable, and maintained, hygiene. This omission was discovered in Q3 and was not addressed even after repeatedly being brought to the attention of ██████████ and ██████████'s appropriate direct report.

Servers (data center)

Of the approximately 500,000 servers in Twitter data centers, ~60% of them are running outdated Operating Systems and, therefore, are non-compliant even with Twitter's own Engineering standards. In addition to security concerns around outdated software components, many of these outdated OSES are not supported by the vendor. They are also not capable of supporting encryption at rest, a critical compliance and Privacy obligations.

¹³ These numbers are rounded due to my direct data access being removed.


Data Center OS Kernel Compliance



While there is a project underway to address Operating System version (kernel) compliance, it has been reportedly long running (multiple years). Progress has not been significant. The project needs to be revisited, revised, re-staffed, and prioritized with clear goals and visible tracking reported to the Risk Committee¹⁴.

On the same engineering dashboards as above it is revealed that ~30% of the software packages running on the ~500,000 data center systems are non-compliant (out of date or need patching).

Data Center Software Compliance



In addition to security and privacy issues, any engineering outage or security event that revealed that the majority of Twitter's production systems are running out of date, and even unsupported, software would likely result in a significant distraction to the

¹⁴ This project appears to have not been appropriately prioritized, and lack appropriate execution plans, while InfoSec [REDACTED].

company. External pressure would be placed upon Twitter to prioritize the addressing of this shortcoming above many #Participation and #Durability efforts.

Both of these situations have been in the same state for the past 12 months.

Processes and Compliance

This topic refers to regulatory obligations and requirements (e.g. SDLC, security reviews, privacy reviews, FTC consent items, regulatory misrepresentations, etc.).

Twitter does not have an industry-appropriate Software Development Life Cycle (SDLC) and Twitter has thus far operated largely without one at all. If it were not for an FTC consent decree, it is possible that Twitter would not be working to put together and deploy an SDLC. This is very atypical in the industry and is a significant risk to the company.

Due to this deficiency, it is inappropriate to label any SDLC progress as “Compliant” to the Committee, as was done in the Information Security documents sent to the Risk Committee. Doing so misrepresents Twitter’s situation as it would be seen by regulators and auditors.

Twitter is in the process of rolling out a registration and SDLC-capable skeleton framework called Flyway. This initial effort, which lacks integration with security reviews, privacy reviews, and other SDLC checks and balances, is unlikely to be viewed as “Compliant” by auditors and investigators.

Making things more challenging, Twitter lacks the ability to provide a count of total software projects (denominator). This means that when someone says a numerator, say “30 projects did privacy reviews”, it’s difficult to know if this is good or bad - is it 30 out of 35 or 3000? The number occasionally being used for context is a count of projects found in the Unified Priority List (UPL). The UPL is a list that at present only represents engineering (i.e. does not include efforts from Site Integrity, Content Moderation, Privacy, InfoSec, IT, Sales, etc.). The UPL further represents only a subset of engineering efforts and does not include day-to-day-running-of-the business software work.

Senior engineers estimate that the UPL represents only a small fraction of the projects that need to adopt a SDLC. To make the number of projects in the UPL even more

problematic as a context value, in 2022 the UPL is intended to change to reflect an even more limited subset of projects than it presently does.


SDLC, security reviews, and privacy reviews, need to be described in terms of whether they would be deemed appropriate by regulators and auditors and in context of how many projects are utilizing these versus how many projects **must** be utilizing them. The UPL is not an appropriate source for total project count.

The Risk Committee members should be aware that *the Twitter SDLC work is not yet what auditors would consider an SDLC, and that security reviews and privacy reviews are not coupled to the SDLC.*

Incidents


Twitter has an unacceptable, and near continuous, number of security and privacy incidents. I estimate there were more than 50 Incidents in 2021¹⁵; approximately an incident per week. Based on my professional experience, peer companies do not have this magnitude or volume of incidents.

H2 2021 had 11 Incidents that were required to be reported to regulators, 5 of which happened in Q4.



The Incidents were predominantly related to areas where Twitter has systemic, long lived, problems: 'Access Control' and 'Security Configuration and Bugs'. Together these problems account for more than 80% of the Incidents.

¹⁵ My notes capture 48 incidents in the period of April - November 2021. 50+ is an extrapolation to include January-March, and December, at the same Incident rate, as my data sources were taken away before this document could be completed.



Meaningfully improving in the 3 areas above (Access Control, Patches and Software Configuration and Versions, and Processes and Compliance) would logically lead the number of incidents to decrease.

Twitter should be experiencing less than 1 regulator-level incident per quarter. Progress in the other areas mentioned, all leading indicators, will drive improvement in the critical lagging indicator of Incidents.

Part 3: Inaccuracies in InfoSec Materials Presented Q4 2021

The Deck


I identified numerous issues in the materials created by [REDACTED] and put forward to the Q4 2021 Risk Committee. I suggested to Mr. Agrawal that I create a corrected replacement deck, including data points. Mr. Agrawal, as CEO, directed that I not create a corrective document and that I send the objectionable deck forward to the Committee. These events are discussed above.

This part of the document examines important inaccuracies and misleading information present in the deck sent to the Q4 Risk Committee.

There were 11 slides in the deck. Two slides were blank. Six slides were qualitative, aspirational, or otherwise did not meaningfully and quantitatively reference critical risks. While there are questions about the appropriateness of these slides for this setting, they are not the focus of this discussion. Three slides included statistics and measurements intended to represent Twitter's environment and key risks. This document now focuses on problems with the data presented in those three slides and articulated at the Committee meeting.

Access Control

Slide 3 and slide 7 inappropriately represented Twitter's access control risk.

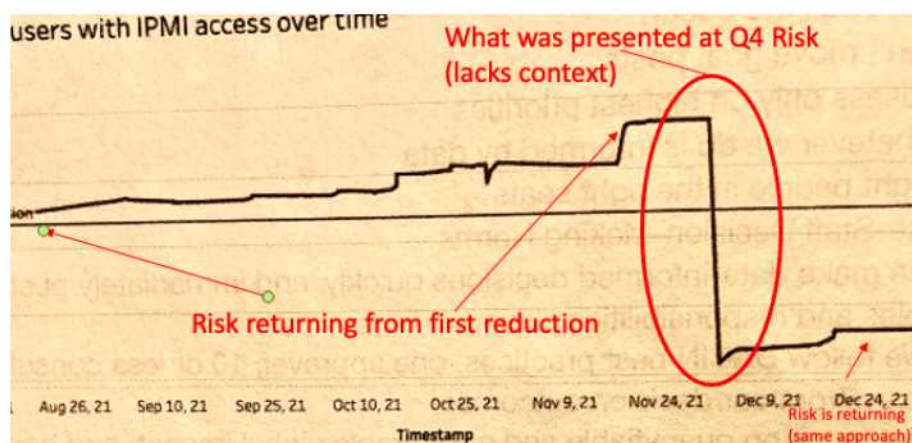


This graphic leads someone to believe significant progress has been made in reducing Twitter access control risk at large. It also implies, with the dotted line, that the reduction is permanent and will continue.

The reduction here focuses on edge cases of access control. These reductions had been made at the end of July and were identified to be temporary in impact, shortly thereafter growing back towards the initial risk.

This graphic represents only 300 users out of the 3,995 users that have production access at Twitter. This represents 5% of the FTEs instead of the 51% needing to be addressed.

The graphic omits that this specific improvement had been previously attempted in July-August. Because there were no systems and solutions in place to enable employees to safely complete their jobs without having the risky access, employees needed to request credentials back to do their work. The Strategy and Operations person overseeing this second reduction, performed again only weeks before the Board and Risk Committee and therefore giving the perception of a recent win, confirmed there were no meaningful changes in the approach this second time that would prevent the re-growth of the risk.



Slide 7 also contains the following statement about access control reduction. The statement is misleading.


Considerations / Next Steps:

Enterprise Identity and Access Management:

- Reduced extraordinary fleetwide Production level access by 66% (up from 34% in Q3)
 - This is part of a much larger effort to reduce access to both fleetwide and partial fleet production access. That effort is focused on reducing the use cases needed for direct production access, and also providing better ways to provide this access in only emergency situations through JIT and other solutions.

The statement above again implies fleetwide Production level access was reduced by 66%. This is not the case. The reduction being described is the subset reduction discussed above. The word “extraordinary” is used to refer to the subset of edge-case FTEs. Without further clarification this is easily misinterpreted and misleading. This is inappropriate to present to the Committee as it stands.

The graphic in the top left area of slide 3, and recreated again bottom left on slide 7, shows a chart intending to reflect the larger issue of broad access to production access to systems and data throughout Twitter FTEs. It is misleading and the data incorrect.



As a reminder, actual fleetwide Production access grew from 46% of FTEs to 51% in 2021. Any Twitter engineer in any country is presently provided direct access to production systems. The accesses to these production systems are not audited.

Next, the dotted line is aspirational and without evidence or existing proofs of working approaches to back it up. In [redacted] oral presentation, [redacted] stated that there was a plan in place to address the larger Access Control risk and that the plan was already

underway. This is incorrect. Engineering, Privacy, and IT (all stakeholders) have expressed significant concern and disagreement over ideas and approaches brought up by InfoSec. InfoSec has not been working collaboratively with stakeholders and there is not an agreed-upon plan¹⁶. The risk is not flat, as portrayed in the graphic above, but rather the risk is meaningfully increasing.

The following image is an accurate depiction of this access control issue at Twitter.



Access to Twitter's Datacenter Production Environment¹⁷

- iii. Dec 2020 46% of employees (2,763 out of 5917)
- iv. Dec 2021 51% of employees (3,995 out of 7714)

(*) The dip was an unintended (internal) incident


SDLC, Security Reviews, Privacy Reviews

Slide 3 top right and Slide 8 - Infographics on Processes and Compliance (e.g. SDLC and Security Reviews)

¹⁶ There is a *goal/s*-type document but it lacks details and it too is contentious among stakeholders for agreement.

¹⁷ It is important to note that this access issue impacts international expansion and operation. In any country where Twitter has an engineer, there is access to production systems and data. Twitter lacks development, staging, and test environments, so they can't mitigate this risk the way other companies do.

SDLC Flyway Total Compliance



This graphic shows adoption of an internal project related to regulatory obligations. It does not show the amount of regulatory *compliance* reached. The term “Compliance” is misleading and inappropriate as a label. The adoption represents a registration-skeleton and is only for a subset of projects at Twitter. See section 2 of this larger document for more details.




The UPL Histogram and Security Review Sparkline are not relevant together

The top right section of slide 3 contains a graphic about security reviews and a histogram of Unified Priority List projects. It is unclear why these two sets of data, Security Reviews and UPL projects, are overlaid.

As [REDACTED] explained when asked what the relationship was between the two datasets during an earlier review of materials, “*there is no correlation*”. The histogram and time series are not connected or related.

This graphic should not have been presented as it can imply a relationship that does not exist. Neither the UPL bars nor the security review line provide appropriate context. They are both numerators without denominators. See section 2 of this document for more details.

The next graphic, security reviews and privacy reviews, is again without context. At best it represents *some amount of work* having been performed but without context. The Y axis only states the number of reviews performed, not how many needed reviews. Similarly there is nothing conveying what was found, themes identified, review targets to hit, or what costs were to the projects and the business.



Security and Privacy Reviews: how many total projects are there and how many are getting security reviews? What are the security reviews finding? Are they worth the cost and how much do they slow down or impact projects and launches?

Patches and Software Configuration and Versions Compliance

Disturbingly absent from the Q4 Risk Committee InfoSec report was information on the state of client and server software compliance, patching, and configuration.

This is a critical area that the FTC has messaged they will evaluate and it is firmly within their domain to do so. As Twitter is in negotiations with the FTC for prior transgressions, where the FTC is now interested in the baseline of security hygiene at Twitter, this is an area of Twitter's the FTC will likely scrutinize. The fact that Twitter is so significantly lacking in this basic security hygiene and practice will contribute to any FTC decisions and enhance penalties.

Only one item in the report presented to the Risk Committee was related to endpoint (employee laptops):

Threat Intelligence:

- Carbon Black Cloud - CBC is now running on 9.7K endpoints
 - CBC replaces the version we were on previously (CarbonBlack Response Standard). CBC has improved endpoint sensor architecture, more stability, and substantially reduces the build time for Twitter engineers. The CBC software is also vastly improved over the previous version, giving us far greater visibility, query performance and capabilities that did not exist in the legacy version.

This entire section is misleading. Stating that Twitter has threat intelligence software running on 9.7k endpoints can sound impressive but is without context. This entire section is misleading. This is the software that has been reporting that the endpoint fleet is extremely out of compliance¹⁸. The prior version of this software, no longer supported by the vendor - hence the upgrade - was reporting the same thing. "Far greater visibility" is without context and, irrespective, it is meaningless if what the software reports is ignored.

Finding that software build times were reduced for a subset of engineers was a happy accident and was not an intentional goal. This brings up an important question: why are engineers performing software builds locally on their laptops. Presently every engineer has a full copy of Twitter's proprietary source code on their laptop. Ideally software builds would be performed on servers in the data centers, or in the cloud, and in an isolated testing environment. The fact that engineers are performing software builds on their laptops (endpoints) and these systems are in such poor security configuration is indeed very disturbing.

The fact that this paragraph of the document does not reference that the security software replaces a previous version of the security software from the same vendor, makes it sound like the change in software was proactive on Twitter's part. It was not. The version of this software already rolled out, on approximately the same number of endpoints listed here, was discontinued by the vendor. Twitter had no choice but to


¹⁸ Across the fleet it is being reported that 30% of laptops have automatic updates disabled, random employee systems have their firewalls turned off, remote desktop turned on, system protection against malicious software disabled, and so on. The software is reporting an endpoint fleet that is in significant disarray.

move to the other software (and quickly). This was not an improvement. This was a lateral move that Twitter was forced to make. Presenting all of this as a win is disingenuous.

A note on Zero Trust and Endpoints (Employee computers)

In regards to endpoint (employee computer) security, [REDACTED] has stated that a “Zero Trust” environment is the go-forward strategy for Twitter. In a Zero Trust environment Twitter employees access internal Twitter services and data without being within a VPN. This approach was popularized by Google around a decade ago and called BeyondCorp. Employee laptops are provided cryptographic credentials and “certified”. The laptops are directly connected to the Internet and (only) the specific service connections into Twitter are encrypted and “protected”. To do this safely requires strong security configurations and hygiene of the laptops receiving Twitter *certification* and ensuring the laptops maintain a strong security posture that is not violated. Moving towards a Zero Trust environment without identifying and addressing the issues with the current state of endpoint configurations and security implies a lack of basic understanding around Zero Trust and information security priorities.

Incidents and Incident Classes




Slide 3 contains information on Incidents: a histogram of Incidents (captured above here) and pie chart of root causes (further below in this document). The histogram is misleading as it only reports a subset of incidents reported and not all SIM (Security Incidents) per quarter. The subset appears to be only related to incidents required to be reported to the Irish Data Protection Commission (a regulatory agency).

It is not appropriate to represent subsets of incidents to the Risk Committee without clearly providing a reason and context. Adding in the missing incidents significantly increases the values. For instance, representing the total number of incidents to the the last section of the chart above, doubles the number of regulator reported Incidents reported to regulators (from 3 to 6) and similarly increases the total Incidents (from 7 to 19). Not only is the chart above incorrect in not representing total incidents, the shape of the chart cannot be trusted to be accurate either.

The correct number is likely closer to 50-60¹⁹ incidents in 2021. More than 1 incident a month was significant and specific enough that it was required to be reported to regulators²⁰. Keep in mind that Twitter is under significant regulatory scrutiny and each of these events worsens Twitter's situation.

The pie chart on Slide 3 showing incident classes (root causes) is also inappropriate.




Pie chart presented to Q4 Risk Committee classifying Incidents

¹⁹ Due to a lack of continued access to data I had to extrapolate to 50-60. I reached this number using the number of incidents from April - November 2021 in my notes and then extrapolated the missing months at the same rate of incidents.

²⁰ Extrapolated from personal notes.

The above assigns less than 7% of incidents with a root cause (class) of access control. This is incorrect. Access Control is the cause of more than 60% of all Incidents, and Security Configurations / Bugs account for almost 25%. These are two areas of critical risk that had not been represented appropriately within Twitter or to the Committee.



Corrected classification of 2021 Incidents - note the correlation of incident classes to critical areas of top risks, as identified in Section 2 of this document. This is a more correct graphic that should have been presented