

AMENDED Exhibit 967

PLAINTIFFS' OMNIBUS OPPOSITION TO DEFENDANTS' MOTIONS FOR SUMMARY JUDGMENT

Case No.: 4:22-md-03047-YGR

MDL No. 3047

In Re: Social Media Adolescent Addiction/Personal Injury Products Liability Litigation

PRIVILEGED & CONFIDENTIAL

2024 Safety Operating Plan ^[E2] | [\[ACP\] PICO/Safety SSP 2023](#)

This OP covers Platform Safety, which includes Trust & Safety Operations (T&S), Law Enforcement Operations (LEO), Safety Engineering, and Safety Product, with valuable input from Legal, Policy, and Communications XFN partners. See [here](#) for more details.

2024 Strategy (Opportunities and Initiatives)

The global safety landscape has grown increasingly challenging. Despite achieving many of our 2023 Safety goals, including launching [Chat Text Reporting](#), [In-App Appeals](#), and [reducing Normalized Violative Views \(NVV\) for CSAM hash matches and Drugs by 75%+](#), serious risks are impacting our community; mounting regulation is hampering our execution; and negative parental perception is impacting our growth. The ongoing fentanyl epidemic in the U.S. and evolving forms of child sexual exploitation (CSE), including financial sextortion, present substantial risk. These issues have contributed to an aggressive regulatory environment and media scrutiny that aim to hold large online platforms, like Snap, accountable. We've needed to react swiftly to new threats (e.g., financial sextortion), platform shortcomings (e.g., NCOSE letter), and emerging compliance requirements (e.g., DSA), diverting us from executing on some of our original 2023 plans. ^[A4] These challenges have also influenced our perception. Compared to our peers, [parents perceive Snapchat as the least safe app](#) and [47% of US teens who don't use Snapchat cite parental restrictions](#) as the primary reason. So, not only is protecting Snapchatters the right thing to do, it is directly connected to the long-term growth of our community and business.

To address the challenges above, we have prioritized the following initiatives:

1. **Continue to bolster our safety foundation and prioritize combating [severe harms](#):** As the preferred communication platform for young people, we have an added responsibility to ensure our Safety foundation protects teens from the most severe abuse (e.g., CSE and lethal drug activity). We will add more upstream safeguards throughout the product, expand our proactive ML detection ^[S8] to more quickly identify abuse, and simplify our in-app reporting tools to make reporting easier and more intuitive, particularly for our youngest users. Operationally, we'll better prioritize teens with quicker SLAs and expand account signal-based enforcement from sextortion to other severe harms. We'll improve NVV and Prevalence measurement to achieve statistical significance more quickly and use the insights to improve proactive defenses to inform our policies and operations.
2. **Proactively identify and mitigate new Safety threats:** Instead of reacting once an issue balloons, ^[J14] we will strengthen our mechanisms and more thoroughly analyze our own data (e.g., report comments, prevalence measurement samples), inputs from trusted sources (e.g., Safety Advisory Board, Law Enforcement agencies, 3P intel vendors), and perform stronger internal adversarial testing to identify and counter emerging risks proactively, earlier in their lifecycle. We'll infuse safety more deeply into Snap's culture through safety-by-design processes to help ensure effective safety defenses are built across the product.
3. **Improve parental perception through more insights, support, and awareness:** We'll continue the momentum we've established with Project Butterfly by focusing on three main areas. First, we'll make Family Center more valuable by providing more insights for parents and highlighting the positive shared interactions parents and teens can have on Snapchat. Second, we'll offer a more compassionate support experience during times of distress for Snapchatters (e.g., Real-Time Support for Imminent/Severe Harm pilot) and their parents (e.g., improved Parental Support flow) and build better tools to support our community directly in Snapchat. Third, because most parents are unaware of our safety protections, we'll raise awareness of Snap's commitment to Safety, locally and nationally, especially in the U.S. ^[15] As a prerequisite, we'll ensure we mitigate the most consequential issues parents care about, like predatory behavior related to sex, drugs, and bullying, ^[J20] and we'll utilize our ongoing research and survey results to track our progress.
4. **Invest in tooling, automation, and vendor accountability to grow costs sublinearly and self-fund new initiatives:** To expand our Safety defenses, we must ensure our operations are scalable. Therefore, we'll invest in solutions to reduce low-quality tasks, automate high-volume queues with low enforcement rates, and improve our agent tooling to reduce AHT. We will also require our vendors to inspect and report the state of their operations more granularly to increase accountability and improve efficiency.
5. **Improve the experience for locked accounts, mitigate over-enforcement, and provide users options for recourse:** As we expand our enforcement efforts, it's even more critical to be fair and equitable in applying our policies. Launching In-App Appeals, In-Tool QA, and the Strike System was a good start; now, we must utilize these new channels as feedback loops between our enforcement precision and policies and ensure we achieve consistently higher accuracy. ^[J23] In 2023, we had to sacrifice the volume and consistency of our FTE QA of vendors' actions ("ground truth QA") to launch Chat Text Reporting and In-App Appeals; in 2024, we will aim for more consistent ground truth QA to ensure focus on quality as we push our vendors to be more productive. We will expand the strike system to take into account more high-fidelity signals across Snapchat, including

Page 1 Comments

- J1 I'm extraordinarily grateful for the work this team has done over the last year. So many critical gaps closed and enormous positive impact for the community. Lots more to do, but really impressed with the progress recently. Thank you.
Jack Brody, 10/4/2023 08:12 PM
- E2 +1 huge thank you to the team.
[REDACTED] 10/6/2023 12:44 PM
- S3 Assuming these "incremental requirements" won't slow down in 2024, do we feel good about our ability to execute on the initiatives listed below?
[REDACTED] 10/1/2023 10:07 PM
- A4 We'll certainly continue to have new regulatory and compliance requirements, but we think the additional investment across key teams being contemplated in the regulatory environment SSP will allow us to execute on these initiatives while also addressing new regulatory requirements.
Aaron Altschuler, 10/2/2023 08:30 AM
- S5 Are we using Gen AI/LLMs to support this work? Per the convo in the eng sec OP, I just want to make sure we aren't handicapping ourselves by relying solely on the existing deals we have with LLM providers for our public-facing use cases.
[REDACTED] 10/1/2023 10:10 PM
- S6 Hi [REDACTED] - if we can remove some of the OpenAI contractual constraints, more things will open up for us. At the same time, there are some privacy considerations as well - we will be sending OAI content, which, today, they store for about 30 days. Perhaps, as part of this, if they can also work with us on reducing the retention time (ideally, they shouldn't have to store it at all), the privacy constraints will also go away. As an example, Google offers us those type of assurances for certain sensitive APIs
- cc: [REDACTED]
[REDACTED] 10/2/2023 12:01 PM
- R7 Hi [REDACTED] -- to be clear, the proactive ML detection referred to here is different from what you and [REDACTED] mentioned.
- These are Snap's ML models (multi-modal, NLP, etc.) that are trained, managed and hosted on Snap infra - we use these models to proactively detect harmful violative content and in some cases automate parts of our trust and safety operations like human review. These are tailor made to capture Snap's policy nuances and performance requirements. Due to efficacy, scale and privacy reasons Safety ML doesn't use 3rd Party ML solutions.
[REDACTED] 10/2/2023 12:39 PM
- S8 Thanks for clarifying, [REDACTED]. Do we use image based models (eg MSFT's X-Ray), or others used by our peers? If so, how's their efficacy?
[REDACTED] 10/6/2023 12:48 PM
- J9 feels like we know when issues are burgeoning before ballooning. How do we know? What would we do. Also seems like we know from say comm's before its in the app? or do we get signals internally that could be even further upstream?
Jerry Hunter, 9/18/2023 04:22 PM
- A10 Thanks in advance for reading this long-ish comment...it's a big topic for our team!
- I would agree that we do tend to find out at the "burgeoning" stage, and the key aspect for our team is ensuring that we *act* before something "balloons". Or at least if we choose not to, it's a conscious choice because we deemed it a lower priority relative to other initiatives.

In terms of identifying issues, there are slightly different types that have some overlap:

-- One type is abuse that's generally emerging across the industry. In this case we'll receive intelligence from third-party trusted sources (Our Safety Advisory Board, various Law Enforcement

Page 1 Comments (Continued)

A10+ action is similar to what we've been doing with financial sextortion

-- The other category is abuse more specific to Snapchat. This is where we will more systematically sample our in-app and support site reports and comments, and prevalence data, while also conducting internal adversarial testing to try and identify issues before anyone else can. We've begun to put this into practice with regular adversarial testing and so far this has helped us identify gaps with moderation, friending, and public profiles. We'll dig into reports and comments more frequently and use that analysis to identify issues early on, so we can be more proactive with our mitigation plans

Alex Osborne, 9/19/2023 12:18 PM

J11 These are good ways. Is it mechanistic or opportunistic?

Jerry Hunter, 9/19/2023 03:31 PM

A12 I think we've gotten more mechanistic this year in some areas - e.g., through proactively gathering and acting on Safety Advisory Board feedback, intel from our 3rd party intel providers, and intentional engagement with LE agencies, parents and educators for this purpose. We've begun other mechanisms (the ones Alex listed in his last paragraph), but we plan to dial up our efforts more in Q4 and next year.

Aaron Altschuler, 9/19/2023 04:09 PM

A13 _Marked as resolved_

Aaron Altschuler, 9/29/2023 08:59 AM

J14 _Re-opened_

Jerry Hunter, 10/6/2023 12:46 PM

15 1 total reaction

reacted with at 2023-10-01 22:11 PM

10/1/2023 10:11 PM

J16 Is this what is driving parental perception? How do we know?

Jerry Hunter, 9/18/2023 04:23 PM

A17 The comms teams and product research team have done great research on this topic recently, and from a combination of surveys and dedicated focus groups, the major issues appear to be related to the sharing and receiving of both solicited and unsolicited explicit content, influenced by our ephemerality and teens being targeted by predators. shared the latest slides via email and here's the link again for reference (<https://drive.google.com/file/d/1AiZoSLsppU-M1R170DCOsDgtC4ZnrJDZ/view?usp=sharing>)

We've also conducted several local outreach efforts, and we often hear about bullying as a major topic that parents and school resource officers are grappling with cc

Beyond the specific types of abuse, perhaps the most notable insight is that parents who use Snapchat are far more likely to think Snapchat is safe for their child to use (see this chart below: <https://docs.google.com/document/d/1J4MIGA3oebw5EmOEcuLzQnqnuMVKTvfKRN77MXCJBu/edit#heading=h.jn9extae7vrz>)

We'll run the sentiment surveys quarterly so we can establish a baseline and track it over time

adding @ and @ in case they want to add more

Alex Osborne, 9/19/2023 12:10 PM

R18 It's noteworthy that drug sales rarely (if ever) gets mentioned in the parent/educator outreach efforts pr the parent focus groups. Though of course we need to continue the drugs work, it doesn't appear to be driving perception with these groups.

Rachel Hochhauser, 9/19/2023 01:39 PM

A19 _Marked as resolved_

Aaron Altschuler, 9/29/2023 09:02 AM

J20 _Re-opened_

J21 what will this look like in practice?

Jack Brody, 10/4/2023 08:09 PM

A22 In practice this is taking our QA data, appeals data and enforcement data, paired with sampling actual cases to understand and correct the gaps in our system and continuously improve the overall customer experience (appropriately removing harmful content and accounts, while not over-enforcing). The gaps could be related to our training (“is our team applying the policies we’ve established correctly?”), the tooling we provide (“does our team have the information/tools needed to make the right enforcement decision?”), the enforcement mechanisms (“is this the right enforcement action for this violation?”) or the actual policies (“is this the right policy given the specific situation?”).

A couple recent actual examples:

— QA/sampling of in app appeals - as we’ve rolled out in app appeals we’ve monitored overturn rates and sampled appeals cases. Many of our appeals are not providing our team enough actionable information. As a result we’re updating the flow to include more validation which will result in more actionable appeals that waste less of our teams time

— QA/sampling of account locks - recently revealed that there is some inconsistency with how we’re applying certain enforcement reasons. This is really important from a customer experience perspective because we have different strike thresholds and device blocking policies depending on the enforcement reason. We’re now planning to use these findings to consider ways to improve our policies, enhance our tooling to give our team more account level signals, update our training based on the policy changes and update the strikes threshold for certain enforcement reasons and update how device blocking is applied

So beyond launching these new products and features we have to make sure we’re constantly tuning them to get the right overall outcomes and this work is that process

That was a very long explanation - please follow up if it didn't answer the question!

Alex Osborne, 10/5/2023 09:44 AM

J23 Very helpful - thank you for the thorough response!

Jack Brody, 10/5/2023 04:17 PM