

# Redesigning Privacy with User Feedback: The Case of Zoom Attendee Attention Tracking

Tony W Li\*  
toli@ucsd.edu

University of California San Diego  
La Jolla, California, USA

Arshia Arya\*  
aarshia@ucsd.edu

University of California San Diego  
La Jolla, California, USA

Haojian Jin  
haojian@ucsd.edu

University of California San Diego  
La Jolla, California, USA

## ABSTRACT

Software engineers' unawareness of user feedback in earlier stages of design contributes to privacy issues in many products. Although extensive research exists on gathering and analyzing user feedback, there is limited understanding of how developers can integrate users' feedback to improve product designs to meet users' privacy expectations. We use Zoom's deprecated attendee attention tracking feature to explore issues with integrating user privacy feedback into software development, presenting public online critiques about this deprecated feature to 18 software engineers in semi-structured interviews and observing how they redesign this feature. Our results suggest that while integrating user feedback for privacy is potentially beneficial, it's also fraught with challenges of polarized design suggestions, confirmation bias, and scope of perceived responsibility.

## CCS CONCEPTS

• **Security and privacy** → *Usability in security and privacy*; • **Software and its engineering** → *Software development techniques*; *Software design techniques*; • **Human-centered computing** → *Empirical studies in HCI*.

## KEYWORDS

privacy by design, privacy engineering, software development, human-centered design

### ACM Reference Format:

Tony W Li, Arshia Arya, and Haojian Jin. 2024. Redesigning Privacy with User Feedback: The Case of Zoom Attendee Attention Tracking. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*, May 11–16, 2024, Honolulu, HI, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3613904.3642594>

## 1 INTRODUCTION

Industry practitioners, in their pursuit of rapid execution, often overlook the broader context of their product design and its potential human and social implications. This developer sentiment is perhaps best exemplified by Facebook's internal motto prior to 2014: "Move fast and break things. The idea is that if you never break anything, you're probably not moving fast enough" [40]. Frequently,

\*First two authors contributed equally to this research.



This work is licensed under a Creative Commons Attribution International 4.0 License.

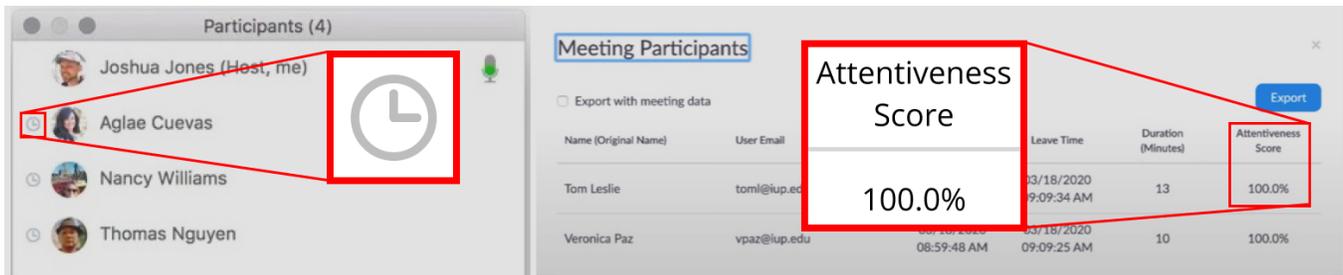
after releasing a new product feature, these companies face significant backlash from users and decide to revoke the feature [9], not only squandering substantial labor but also damaging their reputations.

One recent example is the attendee attention tracking feature in Zoom [51]. Zoom is a video conferencing platform which has seen a huge increase in usage and revenue since the beginning of the COVID-19 pandemic [16] and has rapidly iterated its product to accommodate this growing user base. Zoom developed a feature that allowed the host to monitor the attendees' attention: if Zoom was not the application in focus on a participant's computer for over 30 seconds while someone was sharing their screen, Zoom showed a clock icon next to the participant's name in the participant panel. At the end of each meeting, Zoom also generated a report listing the percentage of time each participant had the presentation window in focus during the meeting (see Figure 1). This feature received significant backlash after launch [48]. The Zoom team later apologized for falling short of the community's privacy and security expectations and decided to remove the attention tracker feature permanently [49].

One important cause for product setbacks like this one is failing to integrate user feedback into the early stages of product design [14, 25, 35]. Earlier research has explored various techniques for gathering user input on privacy aspects of different product designs [6, 17, 25]. These methods range from collecting numerical privacy expectation scores from individual users [25], to identifying privacy norms through the Contextual Integrity privacy framework [6], to classifying privacy concerns extracted from unstructured text [17, 20]. However, a relatively under-investigated question is **how developers might effectively leverage actual user feedback to enhance the privacy of end user products**.

In this paper, we used Zoom attendee attention tracking as a lens to explore the process of integrating user privacy feedback into software development. While there exist different frameworks and approaches to privacy, we focused on the users' privacy expectations [25] in this study, namely user feedback about aspects of a technical feature that contradicted their mental models of how the feature should use their data, as well as how developers utilize those expectations.

We sourced user feedback regarding the Zoom attention tracking feature from public forums, extracting a series of user concerns and organizing them into three categories. We then conducted semi-structured interviews with 18 software engineers of varying seniority at small- to large-scale technology companies. In these interviews, we provided an overview of Zoom's attention tracking feature and its context before asking participants to suggest system



**Figure 1: Attendee attention tracking.** The left screenshot demonstrates the host’s view of inattentive attendees based on whether a participant has Zoom window out of primary focus (marked by gray timer icons in the participants list), and the right screenshot shows an example view of the report generated with the attention tracking (marked by a percentage score for each participant), when this feature was in use. The report view also included participants’ names, emails, join time, leave time, and duration of attendance, as well as attention score. Screenshots adapted from [4, 27].

design changes. We then presented user critiques to the participants, one category at a time, and asked if they would modify their proposed system designs. Throughout the process, we observed how participants processed user feedback and integrated it into their design decisions.

Our results suggest benefits of utilizing user feedback as a checklist for edge cases and as evidence for high-level organizational privacy decisions. However, we also notice redesigning privacy with user feedback is complicated by some challenges. With this format of presenting user feedback, participants tended to devise incremental front-end adjustments which introduced more complicated designs that didn’t scale well for further user concerns, or participants opted to completely abandon the feature in practice; participants deferred on designing solutions for user feedback they personally didn’t agree with or understand; and participants saw some scope of product design, such as different ways users can apply the product, as outside their perceived primary role, limiting engineering solutions to transparently communicating functionality limitations. We discuss future work and recommendations to better understand how software engineers can utilize user feedback to design privacy into software applications, especially in the early stages of design.

## 2 RELATED WORK

Our main objective in this paper is to explore the process of integrating user feedback into privacy-aware software development, which builds on literature from two main areas: (1) collecting users’ privacy feedback and (2) understanding developers’ needs in privacy by design.

**Collecting Users’ Privacy Feedback.** Understanding users’ privacy concerns is crucial for businesses, serving not only as a compliance requirement (e.g., GDPR, CCPA) [45] but also as a foundational step for building trust between companies and users [12]. More broadly, previous work stressed the importance of surfacing user feedback for software evolution [18, 47]. Studies have also investigated users’ privacy concerns in different settings, including IoT [6, 22, 50], online advertising [23, 46], and mobile apps [25, 28, 42], aiming to provide practitioners with a foundational understanding for designing their data practices. For example,

Lin et al. [25] revealed that users became notably concerned about their privacy when informed that the Dictionary app accessed their location. Yet, their concerns significantly diminished when it was clarified that the location data was solely for identifying trending words in their vicinity. Indeed, feedback on users’ privacy has been an important motivator for companies to introduce “purpose strings” into today’s permission system [5].

More recently, researchers have proposed a few ways to collect users’ privacy feedback systematically, tailored for a specific data practice [6, 17]. For example, Apthorpe et al. [6] examined a variety of settings, devices, and data types in specific contexts through questions such as, “A sleep monitor records audio of its owner. How acceptable is it for the monitor to send this information to [different recipients]?” Lean Privacy Review [17] allowed practitioners to gather direct feedback from users for a data practice, in the form of qualitative free-text descriptions and annotated quantitative categories. Feedback collection is especially important as previous work has highlighted differences between software developers’ privacy expectations and those of the users they developed for [25, 37, 39], implying developers struggle to anticipate users’ needs and expectations on their own.

However, user privacy feedback is not self-executing. It requires developers to translate this feedback into actual system design decisions. Our study in this paper focuses on a relatively under-investigated question: *how* do developers use this feedback to enhance the privacy of their products? Specifically, we collected public online critiques about a deprecated feature in popular software: Zoom’s attendee attention tracking. We then presented this feedback to software engineers and observed how they utilized the user feedback to improve the problematic feature’s design.

**Understanding Developers’ Needs in Privacy by Design.** Regulators have embraced “privacy by design” as a critical element of their ongoing revision of current privacy laws [30]. Previous research has explored challenges for integrating privacy requirements into software design [2, 3, 24, 35, 43]. For example, Tahaei et al. conducted interviews with Privacy Champions in software development teams and discovered barriers such as negative privacy culture, internal prioritization conflicts, limited tool support, unclear evaluation metrics, and technical complexity [43]. Other

research has focused on engineer attitudes and practices towards privacy [8, 14, 36], finding engineers face challenges such as lack of perceived responsibility, control, and autonomy; frustrating interactions with legal teams [8]; limiting discourse of privacy to a data security vocabulary; and external organizational climates limiting privacy practice [14].

In contrast, we focus on specific and concrete privacy design tasks, where we instruct developers to integrate privacy into software design and observe their process in order to identify the challenges software engineers face. Our study design was in part inspired by Senarath et al.'s study [35], which asked software engineer participants to design an application for a hypothetical health scenario. During the process, Senarath et al. prompted participants to use multiple privacy frameworks, including Privacy by Design, Data Minimization, Federal Information Processing Standards, and Privacy Impact Assessment, and asked participants to reflect on what they did in the task. Unlike this study, our study is grounded in a real-world data practice and prompts participants with user privacy feedback, exploring privacy from the viewpoint of end users' expectations rather than through a specific framework.

### 3 PILOT STUDY

To inform the design of our main study, we first conducted formative 45-minute semi-structured interviews to understand how software engineers might reflect on a real-world data practice design to improve its privacy centered around anticipated user concerns and privacy expectations.

**Method:** We recruited five software engineers through personal networks for exploratory semi-structured interviews. Inspired by technology probes as a co-designing method [15], we chose the Zoom attendee attention tracking feature as an example to focus design ideas. During the interview, we presented Zoom's attendee attention tracking feature and invited participants to consider different user concerns and ideate on different possible designs of the feature to address those concerns through Crazy Eights exercises [41], a design ideation method commonly used in design sprints where participants are encouraged to come up with multiple ideas in a short period of time (in this case, eight ideas in eight minutes), from which one or more ideas can be further developed into a prototype. All participants had used Zoom as an attendee, host, and presenter in the past, so we also asked participants to informally role-play these different user types to generate more ideas and user concerns from different user types' perspectives. Participants were compensated \$20 USD in the form of shopping gift cards for their participation. The interview results were analyzed in an inductive open coding approach: the two lead authors manually and independently coded up responses then discussed them to agree on a selective coding scheme.

**Results:** This pilot study revealed a couple common participant behaviors. **Participants' ideas focused mainly on 1-2 concerns** that they would encounter as user attendees in their organizations' contexts. Four participants focused their ideas around the same user concern of avoiding false negative measurement of attendees' attention especially with multi-tasking, highlighting the use of multiple screens and passive listening during meetings as common practice. Even after participants role-played as hosts/managers and

presenters as well as attendees, their solutions generally centered around the same attendee concerns. For instance, when considering the manager role, one participant mentioned the same concern they considered as an attendee, saying, *"So at least at my workplace, I know that the people who should be listening would be listening, so I wouldn't care if they [the attendees] are switching windows or they're not paying attention"*.

**Participants' ideas were also high-level and lacked engineering actionability.** Their suggested changes either involved not using the Zoom attention feature at all, or other signals such as *"using 'optional' in calendar to not include people in passive meetings"*, or *"using other interactions such as Q&A during presentations to gauge attentiveness"*. Even when they were asked to role play as other user types, they suggested managers and hosts use broader, more generally "accurate" metrics for tracking participant attentiveness, but were unable to outline clear metric definitions when prompted. Other changes participants suggested added vague incremental front-end polishing to the original user interface; one participant summarized their suggestion to be *"designing the away timer [icon] in a much more sophisticated manner."* Overall, participants generally focused on one or two limited initial concerns they anticipated based on their own experiences as users and devised solutions that lacked specificity and scope for multiple user types, at least in this rapid ideation framework.

Based on these preliminary findings, we observed that engineers could benefit from more guidance, potentially in the form of more diverse user perspectives, to devise more effective and coherent privacy-aware solutions. These results align with previous work suggesting that developers struggle to anticipate user needs and expectations outside of their own experiences [25, 37, 39]. However, we found that participants were able to ideate on a feature such as Zoom attendee attention tracking both as an end user and as a software developer for the product, even if the engineered solutions were a bit limited by their personal user perspective. Thus, we used this pilot study to formulate the methods for our main study, deciding on real user feedback as a way to hopefully ground participants' ideas towards more actionable solutions.

## 4 METHODS

In this section, we describe the recruitment process, protocol, and analysis methods of our study on redesigning privacy with user feedback. We iterated on our pilot study's methodology by introducing user feedback along with the real-world example, sharing the categories of feedback sequentially in random order to generate more ideas to address diverse user concerns. We first gathered and summarized user feedback from publicly available online discussions on the Zoom attendee attention tracking feature, then interviewed software engineers to explore how they would redesign this feature upon reviewing the user feedback we gathered.

### 4.1 Collecting User Feedback by Surveying Public Online Discussions

**Feedback Collection.** The authors surveyed several social media platforms (including Reddit and Twitter) and online publications (including Hackernews [13] and Morning Consult [31]), resulting in 61 historical user comments referencing posts and threads on Zoom

attendee attention tracking. Because this feedback was meant to generate ideas for feature requests for engineers, we only included potentially actionable feedback comments from the search, excluding general complaints and repeat comments from our analysis. Appendix Table 4 enumerates a breakdown of the user feedback sources we used. Zoom attendee attention tracking received more public critique during the onset of the COVID-19 pandemic, so the majority of the collected comments were from the spring of 2020. Given this feature was also deprecated over two years ago, availability of historical user feedback specifically about this feature was somewhat limited, but we did observe saturation with similar feedback before ending this collection. The two lead authors then independently conducted qualitative open coding analyses on these comments and discussed high-level categories for the resulting codes in order to choose representative comments.

**Resulting Codes.** After separately coding the 61 comments, we discovered our prevalent codes aligned with the Belmont Report’s ethical principles [29]. For instance, codes of “Loss of control/autonomy” and “Agency in usage” matched up with the “Respect for Persons” category, while the codes of “Risk of leaking personal data” and “Personal space/data usage” matched up with the “Beneficence” category. After deciding on these ethical principles as the new coding scheme, the authors re-coded the original feedback comments. To decide which comments to show to participants in our study, we selected representative comments that mostly centered around one principle over the others, 3-4 comments for each principle. While there were some positive comments lauding the feature’s usefulness in certain contexts, such as online education, the vast majority of comments were critiques of the feature’s shortcomings. For our representative comments we revealed to study participants, we focused on examples that were more critical, in order to more closely resemble presenting feature requests to our participants. Also, for these representative comments, we did not include minority comments that contradicted the majority comments. Table 1 displays the representative comments selected for each category.

## 4.2 Redesigning Privacy with User Feedback through Semi-Structured Interviews

**Protocol.** We conducted an online 60-minute semi-structured interview with each participant, where we started by asking if they were familiar with Zoom attendee attention tracking as a feature, introducing it if not. In our study, none of the participants were initially aware of its functionality. We clarified that Zoom later removed the feature and this study was meant to explore how this feature could have been different. We first asked participants to redesign the feature from the lens of a user, then introduced the representative comments from each category of user feedback (“Respect for Persons”, “Beneficence”, and “Justice”), one category at a time in a random order, asking the participant to redesign an improved feature addressing the request on top of their previous design (or explain why no change is necessary) (Figure 2). If time allowed, we also encouraged participants to sketch out their ideas, such as in the forms of user storyboards, information flow diagrams, or wireframes, to document how users would utilize the product feature and satisfy their privacy needs (see Figures 3, 4 for examples).

In some cases, sketches were unnecessary, namely if explanations were not feature-specific or did not require illustration, for example if the participant mentioned deprecating the entire functionality. Some participants also preferred to verbally describe their ideas in more detail rather than sketching. After each redesign, the participant was also surveyed on how confidently they believed their current design addressed user needs and concerns, on a 5-point Likert scale. The guide for these semi-structured interviews is included in Appendix A.

**Recruitment.** We recruited 18 software engineers through a snowball sampling strategy [32], where all three authors initially contacted software engineers from their professional and alumni networks then contacted eligible references from the initial participant set. While recruiting participants, we sampled for diversity in years of experience, company size, and privacy engineering experience. 3 participants (17%) identified as female, and the remaining identified as male. Most participants worked in large US-based companies. Participants had an average of 4.3 years of software work experience, and most participants (89%) had not worked with privacy frameworks directly as a collaborator. Two participants reported that their companies had separate security related software and legal teams, and most participants were aware of some approval steps related to privacy before deployment but didn’t actively consider privacy in their work, especially the more junior developers. While most participants had no experience working with privacy frameworks, some had indirect experience (if their product designs had security- or compliance-related considerations) and a couple directly collaborated on or led projects focused on privacy frameworks. Participants’ current employers were based in the United States and India and varied in size, but most participants (94%) worked at medium- (101-1000 people) to large-sized (over 1000 people) companies. Table 2 enumerates a breakdown of this participant information.

**Interview Analysis.** Utilizing recordings and transcripts of the interviews, two authors independently conducted inductive thematic analyses [10] on the data to identify patterns in participants’ responses and grouped those into higher-level descriptions and sub-groups (see Appendix Figure 6 for a general glimpse of our online coding process). We first analyzed the interview recordings to categorize participants’ design suggestions and coded those suggestions into high-level categories (see Table 3). We then dove into understanding the rationale behind each design suggestion and engaged in memo-writing, iterating on the codes to derive analytical themes.

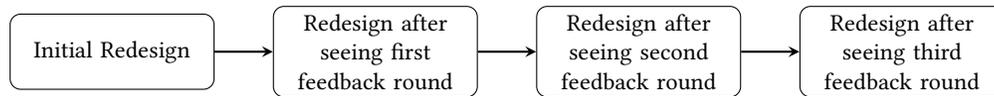
**Ethical Considerations.** Our study was approved by our institution’s IRB. All participants verbally consented to having their data recorded and reported in a scholarly publication. Collected data was stored in a private location accessed only by the authors and anonymized during analysis and reporting. Email information was collected for the sole purpose of scheduling online interviews and distributing \$20 USD compensation in the form of shopping gift cards and was deleted at the completion of the study.

Respect for Persons	Beneficence	Justice
“We wish Zoom would display a notification or let people on the call see whether it’s enabled.”	“That’s kind of intrusive in a way that often doesn’t happen in the physical space or in the physical world. Employers have always had incredible control over how employees spend their time, but the technology makes it faster, more invisible and more sophisticated.”	“Everyone in the meeting room, the host, the speaker, and the attendees should be able to see who’s paying attention.”
“People could be just doing other things and not have the meeting window up. It doesn’t mean they’re not listening.”	“If you have to track people to make sure they pay attention during the meeting, the meeting is pointless and too long. Meetings that are short and packed with useful info nobody wants to miss, are well-attended.”	“I think the issue is not that Zoom knows if its application window has focus, but that it <i>reports</i> focus state to anyone other than the user.”
“Attendees should know how accountability and performance are potentially being ‘graded’ along with the limitations of them.”	“The ‘attention-tracking’ feature [...] was appropriate in some business contexts, but for many new consumers, it presented a privacy conflict.”	“Android users could look at their notifications but not iOS users [...] if you were on Android, you could bring down the Notifications tray for an unlimited amount of time since it didn’t fully cover the Zoom app like the iPhone’s one did.”
“So, my ‘I’m actually watching this on a different screen’ won’t actually fly for much longer?”		“The feature creates a scenario whereby you could be penalized by an employer for doing job-related things, such as checking your notes or updating a memo during an important meeting.”

**Table 1: Categories of user feedback on Zoom attendee attention tracking based on the Belmont Report, along with the example feedback quotes shown to participants in our study.**

#	Total YoE	Experience with Privacy Frameworks	Company Size	Location
P1	6	Indirect experience	Large	US
P2	6	No experience	Large	US
P3	6	No experience	Large	US
P4	3	Indirect experience	Large	India
P5	6	No experience	Large	US
P6	2	No experience	Medium	US
P7	7	No experience	Large	US
P8	9	Indirect experience	Large	US
P9	3	No experience	Large	US
P10	2	No experience	Large	US
P11	6.5	No experience	Large	US
P12	8	Direct experience	Large	US
P13	5	No experience	Large	US
P14	1	No experience	Large	US
P15	2	Direct experience	Medium	India
P16	2	No experience	Small	India
P17	0.5	Indirect experience	Large	US
P18	3	No experience	Large	Spain

**Table 2: Participant information: years of experience (YoE) in software engineering, experience with privacy frameworks, company size, and location.**



**Figure 2: An overview of the interview protocol: after introducing the Zoom attendee attention tracking feature, we asked participants to anticipate concerns as a user and what they would change about the feature, then we showed them the three categories of user feedback in random order, allowing them to update their designs after each feedback reveal.**

## 5 RESULTS

Our results suggest that integrating user feedback into software development for privacy is beneficial but also fraught with challenges. This section discusses the observed benefits and three challenges in detail.

### 5.1 Benefits of User Privacy Feedback

**User feedback as a checklist for edge cases.** User privacy feedback allowed participants to cross-check their assumptions and anticipated concerns. This empowered participants to more confidently and efficiently validate their solutions and make clarifying adjustments in order to address the presented feedback similarly to a checklist. Participants’ reported confidences in their designs also generally increased after seeing user feedback and redesigning further (Figure 5). For example, P5 mentioned that *“being able to compare [my design] against real user feedback made me realize it had accounted for a lot of the original feedback, so then it raised my confidence.”* These responses were to the question, “How confident do you feel this product will satisfy user needs and concerns (on a scale of 1 being not at all confident to 5 being extremely confident)?” The difference in confidence between the distribution of responses on initial designs versus the distribution of responses after the third round of feedback redesigns is also statistically significant on a 95% confidence interval ( $\alpha = 0.05$ ) based on a Wilcoxon signed-rank test ( $p = 0.030$ ).

**User feedback as evidence for high-level organizational privacy decisions.** Participants recognized user feedback as a catalyst for product design changes at a higher organization level they might not normally be able to achieve on their own. As one participant mentioned, *“It does give some more weight to have direct user feedback that goes along with [design decisions]. At least I feel that’s how it operates in [my company]”* (P11). The same participant also mentioned the value of *“some user studies with numbers, because people like numbers. [At my company] it would go a long way to have a more concrete percentage, more than anecdotes.”* Other participants echoed similar sentiments, saying, *“I would say those are pretty valid concerns. But do they really matter? I don’t know. So that’s the reason we want some metrics. We invent some quantitative metrics [...] We should listen to everyone but do the right things. And we need more data to do the right things”* (P12) and, *“In feature development, it’s worth just breaking out these slices of data and just seeing whether there’s a meaningful behavior difference before making decisions”* (P3).

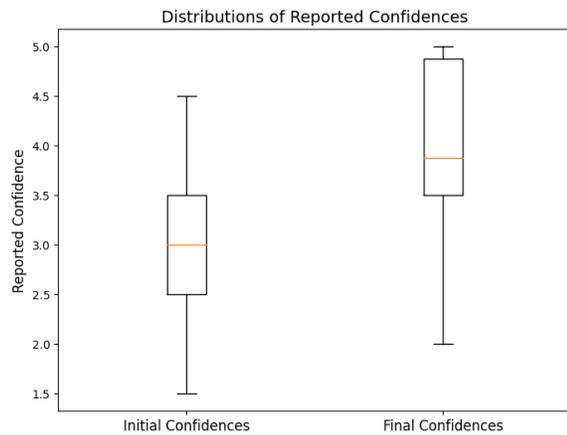
### 5.2 Challenge 1: Polarized Design Suggestions when Redesigning Privacy with User Feedback

At the outset of the study, we anticipated a primary advantage of user privacy feedback to be the ability to enlighten engineers about user concerns and requirements they might not have initially foreseen and enable them to design scalable privacy solutions. After all, previous work identified that acknowledging the differences between developers’ perceived user privacy expectations and those actually expected by users was necessary for developers to successfully integrate privacy into software development [37]. However, we found that user privacy feedback in the form of raw text was unable to empower participants to devise feature updates which could scale for multiple user concerns. We observed a range of different design solutions (Figures 3, 4), but they were mostly focused around either minor, incremental interface changes (e.g. adding notice/consent) or complete deprecation of the feature.

**Incremental interface changes.** When participants saw user feedback they hadn’t originally anticipated, they were able to describe incremental technical solutions, usually focused on the front-end user interface. For example, many of these participants mentioned the best solution was to ensure that attendees understood the data being collected about them and that hosts understood the limitations of the data collected and reported (in a similar vein with purpose strings [12]). These were often in the form of notification interfaces for attendees describing the attention tracking process and “disclaimers” for hosts clarifying suggested usage and best practices for interpreting the reported data (see Figure 3a, 4a). As an example of messaging for hosts, P6 mentioned, *“Instead of saying [attendees are] not paying attention, say they’re not watching their screen, they could still be listening to you.”* In fact, transparency was the most-mentioned category of design suggestions (Table 3), with 16 participants (89%) bringing it up at least once

However, these bandage-like approaches often required many more patches to accommodate other feedback. This buildup of unanticipated changes inadvertently deprioritized usability in some cases. For instance, four participants proposed designs for meeting hosts to customize meeting expectations for different meeting types and different types of attendees (e.g. Figure 3b), but when asked to clarify specifics for the context of attention, a couple mentioned ideas for manually highlighting active times during the meeting to track attention only from certain types of attendee roles determined before the meeting. This type of customizability, while attempting to address a variety of concerns about accuracy in attention tracking, would certainly be a larger overhead for meeting hosts to configure before every meeting they wanted to track.





**Figure 5: Box plot distributions of participants’ reported confidences that their redesigns would satisfy user needs and concerns, both before seeing any user feedback and after seeing all three rounds of feedback. In general, the confidences are higher after seeing feedback.**

Participants who devised these patches for specific concerns also expressed lower confidence and some doubts that their solutions could handle all users’ edge cases; as P14 put it, “*I still can’t completely foresee how it will be used [...] because of that, I’m not gonna give it a 5 [extremely confident].*” P2 even commented, “*I think the more and more I get educated on these corner cases, the more thornier it seems.*” For these participants, it seemed exposure to novel concerns through the feedback humbled their original ideas and forced them to reconsider their privacy design strategies, while still limiting their solutions to patchwork fixes snowballing with front-end interface updates. Addressing concerns on an individual task level introduced issues of scale and usability.

**Feature deprecation.** The challenges of redesigning privacy with user feedback frequently drove participants to consider the opposite extreme solution, feature deprecation, echoing Zoom’s actual solution. Four participants consistently asked if they were allowed to deprecate the feature, or if the company was requiring them to maintain the feature. For instance, these participants asked, “*Am I being pressured to keep this feature?*” (P11) or mentioned “*I don’t think I would implement that feature [...] unless I was forced to*” (P7).

Another significant reason participants struggled to design a solution and considered abandoning the feature was the seemingly unsolvable conflict between their evident moral ground and the technical feasibility. Several participants expressed hesitancy to add more surveillance methods that “*cross the line*” for user privacy but struggled to design technical solutions to track attention as a result. For instance, participants mentioned, “*The idea of surveillance in general tends not to be a great idea. And most people kinda just dislike that idea off the bat*” (P1), “*Despite all the attempts to frame it in a positive way, I don’t think people really like being tracked, just as a human nature kind of thing*” (P5), and “*I think that the act of tracking people inherently makes people suspicious, and I think people are naturally suspicious but have a right to be suspicious, and there will always be some amount of concern about how that data is*

*being used, even if you demonstrate that it’s public and anonymized and aggregated*” (P3). This concern led some participants to believe they couldn’t engineer a technical solution, for instance with some concluding, “*I don’t know how you would be able to resolve this without even more invasive techniques*” (P5), “*I don’t think you’ll be able to capture [attention] digitally in any sort of digital platform or environment*” (P1), and “*I think anything that I would propose would be too invasive of privacy, and I would not be comfortable implementing that*” (P7). It’s worth noting that several participants seemed to assume an inherent organizational or business pressure to keep the product feature. As P6 mentioned, “*The fact that I built this feature in the first place meant that some other customer of mine asked for it, and if I don’t give them this feature, then they might just go somewhere else.*”

Upon closer inspection of the participants’ reasons to deprecate the product, participants also seemed to pull in personal opinions and experiences outside of the displayed user feedback. Before seeing any feedback, P11 even noted, “*I don’t think using it is good [...] it should just not exist. Because anything I can think of to, like, improve the quality of data is just more recording and analyzing. And they’re all very subjective to whoever creates the logic [...] so I think there’s a lot of human factor and decision, and I don’t think any of it is reliable and should not be presented to the end user.*” As another example, P7 justified deprecating after seeing all three rounds of user feedback but for reasons beyond any specific feedback, mentioning, “*I don’t think it effectively tracks performance, and I think it erodes trust and culture within the company. And it would cause micromanaging within the teams. And I think people would try to find ways to gamify the system [...] and I feel like there’d be more politicking involved. I think this would not be good for company culture. Yeah, I foresee it as like a detriment [...] I think it shuts down voices potentially.*”

### 5.3 Challenge 2: Confirmation Bias in Integrating User Privacy Feedback

We had also anticipated user privacy feedback to foster empathy between developers and the end users, but we found that this format of raw user privacy feedback was limited by some apparent confirmation bias. Section 5.1 describes how feedback that aligned with participants’ expectations was effective for boosting participants’ confidences on their designs, but when feedback ran counter to participants’ expectations, participants were much more hesitant to suggest product fixes to address those concerns.

Some of these participants deferred evaluation of these concerns to others, such as product managers, to decide on task prioritization for engineers, while a couple of participants plainly refuted the value of addressing those concerns. For instance, for one user comment they disagreed with, P2 decided, “*I think I would essentially pass the buck to like PM’s or UX researchers to essentially determine, is that a worthwhile business problem to solve? Like how much value does that generate for our users, whether it’s like they feel safer.*” P16 refuted one feedback comment, mentioning, “*That just makes this feature redundant. I don’t think this should be addressed [...] No one cares about that. So if I’m a developer at Zoom, I wouldn’t want to implement this feature, because there’s no point.*” In general, participants who did not connect with user feedback seemed less

Design Suggestions	Examples of Design Suggestions in this category
Business Needs	Survey preferences from premium customers; Defer to other stakeholders
Consent	Opt in; Opt out
Context Dependence	Collect attention data differently for different meeting types or user roles
Data Retention Policy	Anonymize/aggregate attention data; Store raw data locally and only send metadata to servers
Data Usage Policy	Show data analytics to optimize meetings rather than evaluate individual performance
Deprecation	Remove the feature overall
Feature Definition	Alter time-based criteria (e.g. 5min instead of 30sec); Expand the definition of attention
Surveillance signals	Add other engagement metrics to track attention; Avoid more invasive technologies
Transparency	Notify when tracking; Document and communicate the feature clearly

**Table 3: Design suggestion categories and examples of exact design suggestions made by the participants. This table encapsulates 30 unique design suggestions, broadly categorized into these nine areas.**

motivated to take the initiative to fix them or looked to others to verify if they were worth addressing.

### 5.4 Challenge 3: Perceived Responsibility in Designing for Privacy

Another common challenge we observed was participants viewing product decision-making as outside of their primary role or acknowledging that outside factors could thwart their efforts to make a difference in their organization, even when presented with user feedback.

**Engineering solutions accounting for different user applications of the product were limited to clear communications of feature functionality.** There was a common pattern of participants considering some extent of product usage as outside the scope of their software design work. For instance, when considering how this tool would be used by employers for tracking attention of their employees in work meetings, many participants acknowledged how punitive action based on this tool could be seen as unfair and could even be abused to “*make arguments in bad faith*” (P3), but most of those participants mentioned this consideration as out of scope of the product design or struggled to think of ways to address it, other than to add clarity, such as with the previously-mentioned “*disclaimers*” about the data for users. For instance, participants mentioned, “*You can’t guarantee that companies aren’t gonna [...] use this as justification for axing hybrid or fully remote work, right? Because that’s beyond a Zoom product designer’s control*” (P3), “*We provide the information, but [...] it’s really up to the meeting user, whatever the context of the meeting is [...] but as for what it means to be accountable, I don’t think the tool needs to answer that*” (P1), and “*Some things you just can’t control, like which environment [...] and what the employer is using this data for [...] that is the responsibility of the employer and the employee*” (P4). Others echoed similar sentiments but mentioned clear communication, both from the product and from intermediate users such as organization heads and meeting hosts, as the best solution for this case, saying, “*I don’t think we can address sort of, like, how it should be used. But we can provide enough information and enough clarity on how much information we do provide for users of it to make sort of those decisions kinda outside of the feature*” (P1), or even, “*Do companies have a right to use this information against you when they’re evaluating a performance? So*

*I think the company should be very clear about this, only then the feature should be enabled*” (P16). Other participants stated potential misuse as a factor to simply deprecate the feature, also struggling to think of other design solutions. For instance, P8 mentioned, “*We built this tool that the employer may be trying to force on people, and it’s more on the employer for using it a certain way [...] But that raises the question of if this feature should even exist, if it can cause these sorts of controversies.*”

In general, these participants did not consider downstream implications for end users as central for their engineering role in product design, though some of them advocated for transparently communicating how the data can be used so that users could address those implications, while others considered removing the feature entirely.

**Participants deferred to others to determine how the broader organization balanced feature prioritization with user privacy comfort.** A number of participants mentioned a consideration of realistic business factors, namely that the product would likely cater to the highest-paying customers. P3 recognized that for many product decisions, “*It depends on who your highest paying customer is and what they want,*” and P13 similarly noted, “*I think it really ultimately depends on whether or not that’s a feature that the paying customers want to put in, ‘cause if you end up putting that in as a feature change, and the people that actually pay for the licenses don’t like it, then there’s no point actually putting it in.*” These participants assumed that they as engineers would concede to what the company decided would be best for their paying customers, with a couple participants explicitly mentioning they would prioritize surveying preferences from premium users (or potential paying users) and defer to them to decide on product direction. In this sense, not only did engineers recognize their employer and the end user as stakeholders in the design process, but also the perceived relative importance (in this case, financial weight) of certain users as affecting the product design decisions.

Several other participants also cited external regulations as taking precedence when guiding product decisions. For example, P8 noted, “*I would assume that Zoom has some sort of security council and legal counsel [...] to discuss this type of point, because there’s also different countries – I know the European Union has different security laws regarding user data,*” and P12 mentioned, “*First, we have to make sure we follow the laws, regulations – GDPR, some common*

*legal codes, something like that.*” In general, these broader balance considerations of employer and valued customer priorities as well as external regulations were seen as out of scope for these participants in the role of an engineer.

## 6 LIMITATIONS

This paper uses Zoom attendee attention tracking as the lens to explore the process of redesigning privacy with user feedback. We chose this feature since many software engineers are familiar with the software and the feature’s context is relatively simple. This feature also comprises a complete data pipeline, covering notice and consent, data processing, and data storage and retention. However, users might be biased by the brand name [34] and their specific work environments, as well as knowledge of its real-world deprecation. Further research may explore redesigning privacy with more diverse and potentially anonymized use cases.

In addition, we acknowledge that completely avoiding bias from recruiting is nearly impossible, especially with our snowball recruiting methods. We limited our study’s scope to one of exploratory ideation rather than a full evaluation of designs. We encourage future studies to focus on more diverse participant demographics, perhaps with alternative recruitment methods with sufficient pre-screening [11, 19, 44] and best practices inspired by research in other areas such as security [38]. Also, because our study was more exploratory, we conducted 45-60 minute interviews; however, alternative methods such as focus groups or in-depth ethnographic observations, as well as evaluative research on design prototypes, may lead to more detailed results.

This study also focused on privacy expectations as expressed by users and thus limited its approach to privacy as informed by user feedback; future work could explore different theoretical and legal approaches to privacy by design and how to better inform new legislation to guide software development processes, especially in different countries.

## 7 DISCUSSION

This study aims to inform the design of a system that can help developers integrate users’ feedback to improve product design to meet users’ privacy expectations. Our findings indicate that presenting user privacy concerns as raw text feedback to software engineers has limited effectiveness for several reasons. Here we discuss potential solutions and research directions based on our observed challenges.

**Addressing polarized designs.** In our study, we provided developers with raw feedback text categorized according to ethical principles [29]. However, addressing this varied feedback might necessitate different degrees of modification to the system. In our study, we observed this in participants’ polarized design decisions, which did not always account for inter-dependencies in design and were even contradictory. Only presenting raw feedback text makes it hard for developers to prioritize their fixes as well. Future systems should consider developing a more formalized and standardized protocol for presenting feedback to software engineers, potentially

by offering the feedback in a more organized and consolidated format.

**Addressing confirmation bias.** Raw feedback batched in categories may have exposed participants to unexpected or new user concerns, but it could not foster sufficient empathy between developers and end users to overcome developer confirmation bias. This study is not the first to observe personal opinions affecting engineers’ privacy work [7, 35], so further research should explore better ways to enable this empathy. While user feedback in this study seemed effective as a confirmatory medium through which to connect developers with users, more methods could be explored beyond simply surfacing raw feedback, such as incorporating role-play as suggested in prior work [37], or using feedback to expose a product’s usefulness and practical results, which have been shown to affect engineers’ intentions to follow privacy engineering methodologies [36].

**Addressing perceived responsibility.** In our study, we observed that user feedback alone was insufficient for engineers to autonomously make high-level privacy decisions, but participants mentioned that feedback and quantitative data can inform decisions made by others in their organizations. Data as evidence provides more weight behind decisions and could garner more trust in organizations.

For the participants, privacy was a collaborative effort that required organizational buy-in, as there seemed to be an inherent and sometimes-explicit tension between end user concerns and organizational pressures to maintain an existing product that served business customers. Participants implied broader product decisions were usually organizational- or management-driven, and engineers didn’t have much influence to reconcile their personal viewpoints with the broader organizational goals, but user feedback and, perhaps more strongly, quantitative usage data would empower them to bring up user concerns with management. Future work to prepare engineers for that conversation could empower engineers to promote human-centered privacy design.

## 8 FUTURE WORK

Future work is needed to investigate how to overcome the challenges we observed and potentially take advantage of and expand the benefits we identified. Here we specify relevant research communities and outline other potential future directions based on this study.

**Research communities.** This area would benefit from further mutual understanding and collaboration between human-computer interaction (HCI) and software engineering communities. For instance, HCI researchers can develop ways of gathering and analyzing user feedback to improve software design, but we found that those efforts will be limited if they don’t account for software process realities, such as perceived responsibility. Similarly, this study can inform software researchers on how to utilize user feedback in software engineering processes, such as with a metrics-based evidence pipeline for engineers to inform higher-level product decisions rather than second-guess their individual role capacity for

product changes. This community collaboration is crucial for future work in order to better integrate user feedback into practical software processes, especially for iterative privacy design.

**User feedback beyond critiques.** Our study focused on constructive critiques and formulated them as feature requests. However, other types of user feedback, such as positive feedback, could also be explored. For example, instead of completely focusing on designs addressing user complaints, noting aspects praised by users in other successful products and transferring them to new product designs is a reasonable approach in product design; indeed, this is a component of competitive analysis in fields such as business [1] and user experience design [33]. Exploring how to surface those positive comments in the context of software development could prove fruitful.

In addition, this study surfaced feedback through social media. Other sources and more direct feedback loops, such as targeted in-app surveys, could add insight into the effect of different types of feedback on software development. With recent advancements in generative artificial intelligence, it could also be worthwhile to explore ways of generating simulated user feedback in formats beneficial for product designers.

**Scaling real-world user feedback.** Building on previous work [21, 35], our study could potentially be extended to investigate how to form privacy guidelines, especially after identifying challenges engineers faced with designing for privacy from unstructured user feedback. Historical data on user feedback to inform how privacy guidelines should shift over time, and how to surface that data, could also be informative. Research to understand how user feedback can inform privacy guidelines, such as with identifying loopholes of existing regulations, can bring privacy practice closer to user privacy expectations. This could be achieved systematically by crawling online critiques at scale and analysing them qualitatively with a bottom-up approach.

## 9 CONCLUSION

This research explores the concept of redesigning privacy with user feedback, empirically observing industry practitioners in interviews and analyzing their redesign ideas, behaviors, and thought processes. Our results indicate user feedback shows promising benefit as a checklist for engineering edge cases and as justification for high-level organizational privacy decisions. However, redesigning privacy with user feedback suffers from 3 main challenges: polarized design suggestions, confirmation bias, and limits of perceived responsibility. Future work to address these challenges while taking advantage of user feedback benefits, as well as expanding and validating these results in broader contexts, would be beneficial for understanding how to better integrate user privacy into product design.

## REFERENCES

- [1] U.S. Small Business Administration. 2023. Market research and competitive analysis. <https://www.sba.gov/business-guide/plan-your-business/market-research-competitive-analysis>. (Accessed on 12/12/2023).
- [2] Yaqoob Al-Slais. 2020. Privacy Engineering Methodologies: A survey. In *2020 International Conference on Innovation and Intelligence for Informatics, Computing and Technologies (3ICT)*. 1–6. <https://doi.org/10.1109/3ICT51146.2020.9311949>
- [3] Abdulrahman Alhazmi and Nalin Asanka Gamagedara Arachchilage. 2021. I'm all ears! Listening to software developers on putting GDPR principles into software development practice. *Personal and Ubiquitous Computing* 25, 5 (May 2021), 879–892. <https://doi.org/10.1007/s00779-021-01544-1>
- [4] Jenna Amatulli. 2020. Zoom Can Track Who's Not Paying Attention In Your Video Call. Here's How. | HuffPost Life. [https://www.huffpost.com/entry/zoom-tracks-not-paying-attention-video-call\\_1\\_5e7b96b5c5b6b7d80959ea96](https://www.huffpost.com/entry/zoom-tracks-not-paying-attention-video-call_1_5e7b96b5c5b6b7d80959ea96). (Accessed on 07/26/2023).
- [5] Apple Developer. 2023. Write clear purpose strings - Tech Talks - Videos - Apple Developer. <https://developer.apple.com/videos/play/tech-talks/110152/>. (Accessed on 09/13/2023).
- [6] Noah Apthorpe, Yan Shvartzshnaider, Arunesh Mathur, Dillon Reisman, and Nick Feamster. 2018. Discovering smart home internet of things privacy norms using contextual integrity. *Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies* 2, 2 (2018), 1–23.
- [7] Oshrat Ayalon, Eran Toch, Irit Hadar, and Michael Birnhack. 2017. How Developers Make Design Decisions about Users' Privacy: The Place of Professional Communities and Organizational Climate. In *Companion of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (Portland, Oregon, USA) (CSCW '17 Companion)*. Association for Computing Machinery, New York, NY, USA, 135–138. <https://doi.org/10.1145/3022198.3026326>
- [8] Kathrin Bednar, Sarah Spiekermann, and Marc Langheinrich. 2019. Engineering Privacy by Design: Are engineers ready to live up to the challenge? *The Information Society* 35, 3 (2019), 122–142. <https://doi.org/10.1080/01972243.2019.1583296> arXiv:<https://doi.org/10.1080/01972243.2019.1583296>
- [9] Chris Benner and Kung Feng. 2020. Elon Musk reflects Silicon Valley's 'move fast and break things' culture. *San Francisco Chronicle* (2020). <https://www.sfchronicle.com/opinion/openforum/article/Elon-Musk-reflects-Silicon-Valley-s-move-15271652.php>
- [10] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3 (01 2006), 77–101. <https://doi.org/10.1191/1478088706qp063oa>
- [11] Anastasia Danilova, Alena Naiakshina, Stefan Horstmann, and Matthew Smith. 2021. Do You Really Code? Designing and Evaluating Screening Questions for Online Surveys with Programmers. In *Proceedings of the 43rd International Conference on Software Engineering (Madrid, Spain) (ICSE '21)*. IEEE Press, 537–548. <https://doi.org/10.1109/ICSE43902.2021.00057>
- [12] Müge Fazlıoğlu. 2023. IAPP Privacy and Consumer Trust Report – Executive Summary. <https://iapp.org/resources/article/privacy-and-consumer-trust-summary/>. (Accessed on 09/13/2023).
- [13] Hacker News User. 2020. [dupe] Using Zoom? Here are the privacy issues you need to be aware of. <https://news.ycombinator.com/item?id=22664219>. (Accessed on 08/07/2023).
- [14] Irit Hadar, Tomer Hasson, Oshrat Ayalon, Eran Toch, Michael Birnhack, Sofia Sherman, and Arod Balissa. 2018. Privacy by designers: software developers' privacy mindset. *Empirical Software Engineering* 23 (2018), 259–289.
- [15] Hilary Hutchinson, Wendy Mackay, Bo Westerlund, Benjamin B Bederson, Alison Druin, Catherine Plaisant, Michel Beaudouin-Lafon, Stéphane Conversy, Helen Evans, Heiko Hansen, et al. 2003. Technology probes: inspiring design for and with families. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 17–24.
- [16] Mansoor Iqbal. 2023. Zoom Revenue and Usage Statistics (2023). *Business of Apps* (2023). <https://www.businessofapps.com/data/zoom-statistics/>
- [17] Haojian Jin, Hong Shen, Mayank Jain, Swarun Kumar, and Jason I Hong. 2021. Lean privacy review: Collecting users' privacy concerns of data practices at a low cost. *ACM Transactions on Computer-Human Interaction (TOCHI)* 28, 5 (2021), 1–55.
- [18] Jan Ole Johanssen, Anja Kleebaum, Bernd Bruegge, and Barbara Paech. 2019. How do Practitioners Capture and Utilize User Feedback During Continuous Software Engineering?. In *2019 IEEE 27th International Requirements Engineering Conference (RE)*. 153–164. <https://doi.org/10.1109/RE.2019.00026>
- [19] Harjot Kaur, Sabrina Amft, Daniel Votipka, Yasemin Acar, and Sascha Fahl. 2022. Where to Recruit for Security Development Studies: Comparing Six Software Developer Samples. In *31st USENIX Security Symposium (USENIX Security 22)*. USENIX Association, Boston, MA, 4041–4058. <https://www.usenix.org/conference/usenixsecurity22/presentation/kaur>
- [20] Katharina Krombholz, Adrian Dabrowski, Matthew Smith, and Edgar Weippl. 2017. Exploring design directions for wearable privacy. (2017).
- [21] Marc Langheinrich. 2001. Privacy by design—principles of privacy-aware ubiquitous systems. In *International conference on ubiquitous computing*. Springer, 273–291.
- [22] Josephine Lau, Benjamin Zimmerman, and Florian Schaub. 2018. Alexa, are you listening? Privacy perceptions, concerns and privacy-seeking behaviors with smart speakers. *Proceedings of the ACM on human-computer interaction* 2, CSCW (2018), 1–31.
- [23] Hao-Ping Hank Lee, Jacob Logas, Stephanie Yang, Zhouyu Li, Nata Barbosa, Yang Wang, and Sauvik Das. 2023. When and Why Do People Want Ad Targeting Explanations? Evidence from a Four-Week, Mixed-Methods Field Study. In *2023 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2903–2920.

- [24] Tianshi Li, Elizabeth Louie, Laura Dabbish, and Jason I Hong. 2021. How developers talk about personal data and what it means for user privacy: A case study of a developer forum on reddit. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW3 (2021), 1–28.
- [25] Jialiu Lin, Shahriyar Amini, Jason I Hong, Norman Sadeh, Janne Lindqvist, and Joy Zhang. 2012. Expectation and purpose: understanding users' mental models of mobile app privacy through crowdsourcing. In *Proceedings of the 2012 ACM conference on ubiquitous computing*. 501–510.
- [26] Miro. 2023. Miro. <https://miro.com/>. (Accessed on 09/14/2023).
- [27] Dr. Veronica Paz. 2020. Zoom Attendance Tracking. <https://www.youtube.com/watch?v=01lbmOWFRc8>. (Accessed on 09/13/2023).
- [28] Ashwini Rao, Florian Schaub, Norman Sadeh, Alessandro Acquisti, and Ruogu Kang. 2016. Expecting the unexpected: Understanding mismatched privacy expectations online. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. 77–96.
- [29] Michelle Roth-Cline and Robert Nelson. 2014. The ethical principle of scientific necessity in pediatric research. 2014. The Belmont report. Ethical principles and guidelines for the protection of human subjects of research. *The American Journal of Bioethics* 14, 12 (2014), 14–15.
- [30] Ira Rubinstein and Nathaniel Good. 2012. Privacy by Design: A Counterfactual Analysis of Google and Facebook Privacy Incidents. *SSRN Electronic Journal* 28 (08 2012). <https://doi.org/10.2139/ssrn.2128146>
- [31] Sam Sabin. 2020. How to run a Crazy eights workshop. *Prototypr* (2020). <https://pro.morningconsult.com/articles/remotework-video-conference-data-privacy>
- [32] Georgia Robins Sadler, Hau-Chen Lee, Rod Seung-Hwan Lim, and Judith Fullerton. 2010. Recruitment of hard-to-reach population subgroups via adaptations of the snowball sampling strategy. *Nursing & health sciences* 12, 3 (2010), 369–374.
- [33] Amy Schade. 2013. Competitive Usability Evaluations: Learning from Your Competition. *Nielsen Norman Group* (2013). <https://www.nngroup.com/articles/competitive-usability-evaluations/>
- [34] Stuart Schechter and Cristian Bravo-Lillo. 2014. Using ethical-response surveys to identify sources of disapproval and concern with Facebook's emotional contagion experiment and other controversial studies. (2014).
- [35] Awanthika Senarath and Nalin AG Arachchilage. 2018. Why developers cannot embed privacy into software systems? An empirical investigation. In *Proceedings of the 22nd International Conference on Evaluation and Assessment in Software Engineering 2018*. 211–216.
- [36] Awanthika Senarath, Marthie Grobler, and Nalin Asanka Gamagedara Arachchilage. 2019. Will They Use It or Not? Investigating Software Developers' Intention to Follow Privacy Engineering Methodologies. *ACM Trans. Priv. Secur.* 22, 4, Article 23 (nov 2019), 30 pages. <https://doi.org/10.1145/3364224>
- [37] Awanthika R. Senarath and Nalin Asanka Gamagedara Arachchilage. 2018. Understanding user privacy expectations: A software developer's perspective. *Telematics and Informatics* 35, 7 (Oct. 2018), 1845–1862. <https://doi.org/10.1016/j.tele.2018.05.012>
- [38] Raphael Serafini, Marco Gutfleisch, Stefan Albert Horstmann, and Alena Naiakshina. 2023. On the Recruitment of Company Developers for Security Studies: Results from a Qualitative Interview Study. In *Nineteenth Symposium on Usable Privacy and Security (SOUPS 2023)*. USENIX Association, Anaheim, CA, 321–340. <https://www.usenix.org/conference/soups2023/presentation/serafini>
- [39] Swapneel Sheth, Gail Kaiser, and Walid Maalej. 2014. Us and Them: A Study of Privacy Requirements across North America, Asia, and Europe. In *Proceedings of the 36th International Conference on Software Engineering (Hyderabad, India) (ICSE 2014)*. Association for Computing Machinery, New York, NY, USA, 859–870. <https://doi.org/10.1145/2568225.2568244>
- [40] Wired Staff. 2012. Mark Zuckerberg's Letter to Investors: 'The Hacker Way'. <https://www.wired.com/2012/02/zuck-letter/>
- [41] Hana Stevenson. 2019. How to run a Crazy eights workshop. *Prototypr* (2019). <https://blog.prototypr.io/how-to-run-a-crazy-eights-workshop-60d0a67b29a>
- [42] Mohammad Tahaei, Ruba Abu-Salma, and Awais Rashid. 2023. Stuck in the Permissions With You: Developer & End-User Perspectives on App Permissions & Their Privacy Ramifications. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–24.
- [43] Mohammad Tahaei, Alisa Frik, and Kami Vaniea. 2021. Privacy champions in software teams: Understanding their motivations, strategies, and challenges. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [44] Mohammad Tahaei and Kami Vaniea. 2022. Recruiting Participants With Programming Skills: A Comparison of Four Crowdsourcing Platforms and a CS Student Mailing List. In *CHI Conference on Human Factors in Computing Systems (CHI '22)*. ACM. <https://doi.org/10.1145/3491102.3501957>
- [45] TechCrunch. 2023. Privacy goes global: Evolving consumer expectations and the law | TechCrunch. <https://techcrunch.com/sponsor/usercentrics/privacy-goes-global-evolving-consumer-expectations-and-the-law/>. (Accessed on 09/13/2023).
- [46] Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. 2012. Smart, useful, scary, creepy: perceptions of online behavioral advertising. In *proceedings of the eighth symposium on usable privacy and security*. 1–15.



**Figure 6: Screenshot of two authors' thematic analysis process, showing interview transcript quotes grouped under different codes. This analysis was conducted in Miro [26].**

- [47] Simon van Oordt and Emitza Guzman. 2021. On the Role of User Feedback in Software Evolution: a Practitioners' Perspective. In *2021 IEEE 29th International Requirements Engineering Conference (RE)*. 221–232. <https://doi.org/10.1109/RE51729.2021.00027>
- [48] Tom Warren. 2020. Zoom faces a privacy and security backlash as it surges in popularity. *The Verge* (2020). <https://www.theverge.com/2020/4/1/21202584/zoom-security-privacy-issues-video-conferencing-software-coronavirus-demand-response>
- [49] Eric Yuan. 2023. A Message to Our Users | Zoom Blog. <https://blog.zoom.us/a-message-to-our-users/>. (Accessed on 07/26/2023).
- [50] Serena Zheng, Noah Aporthe, Marshini Chetty, and Nick Feamster. 2018. User perceptions of smart home IoT privacy. *Proceedings of the ACM on human-computer interaction* 2, CSCW (2018), 1–20.
- [51] Zoom. 2023. Attendee attention tracking – Zoom Support. <https://support.zoom.us/hc/en-us/articles/115000538083-Attendee-attention-tracking>. (Accessed on 07/26/2023).

## A FIGURES AND TABLES APPENDIX

Table 4 details distribution of sources for the potentially-constructive comments concerning Zoom's attendee attention tracking feature. Figure 6 displays a high-level overview of the thematic analysis process conducted in Miro, where participant quotes (transcribed to notes) were grouped into open-coded categories and subcategories, such as "Collect more data/info before change" and "Kill it".

Feedback Source	Count
Twitter	17
Reddit	12
Hackernews	12
Youtube	4
Morning Consult	3
Gadget Hacks	2
Vice	2
Our pilot study	2
Miscellaneous blogs/articles	7

**Table 4: Sources of user feedback and the counts of comments analyzed from each source. Most were sourced from Twitter, Reddit, and Hackernews.**

### A.1 Interview Guide

The following is the guide agreed on by the two interviewers for the semi-structured interviews for this paper's main study. In the spirit of semi-structure, individual interviews varied in interviewer questions and follow-up probes, but the general interview structure followed this guide.

**Intro [1 min]**

Hello! My name is \_\_\_\_ and I'm doing an interview about workplace surveillance as part of my studies at my institution. Thank you for agreeing to take the time to talk today. This interview will take 45-50 minutes of your time. It is also voluntary, so you may choose to withdraw from the interview at any point for any reason. And please don't hesitate to ask questions at any point.

**Consent [1 min]**

For the purposes of my study, I will be recording our conversation as well as any sketches I may ask you to do during this interview. I will anonymize your identity and will report the results only in the context of academic publications. Do I have your consent to use your data, and would you still like to participate in this study?

**Warmup [3 min]**

Thank you! To get things rolling, let's start by going through a few questions:

- What is your current occupation?
- [If not currently SWE or in tech] Have you worked as a software engineer at a company within the last 10 years? Think back to when you did work as one.
- If you're comfortable sharing, may I ask what company you work at?
- [If not] Would you be okay with sharing the approximate size of your company?
- How long have you worked at your company?
- Approximately how large is your direct team?
- Have you ever worked with any privacy frameworks (such as Fair Info Practices, Privacy by Design, or Data Minimization) in your software designs?
- [If not] How familiar are you with privacy frameworks?

**Initial Brainstorm [10 min]**

Have you heard of Zoom's attendee attention tracking feature? It's a feature on the Zoom video calling platform that isn't active anymore, but basically if someone was sharing their screen in a group call, the hosts could see which participants weren't actively on the presentation window for more than 30 seconds. Inactivity was marked by a gray timer icon next to their name on the participants list (visible only to the hosts). After the meeting, Zoom would generate a report listing the percentage of time each participant had Zoom in focus during the meeting, as well as how long they were in the meeting. [Show screenshots.]

I'd like you to first imagine you were a user in a meeting utilizing Zoom Attention. Can you anticipate any needs and concerns as a user?

How might these issues be addressed in a feature update?

How confident do you feel this product will satisfy user needs and concerns (on a scale of 1 being not at all confident to 5 being extremely confident)?

- [Based on the ideas, ask for follow-up explanation, such as with an info flow, user journey storyboard, privacy storyboard, system architecture diagram, etc.]

- [Also consider probing about what data is collected, data retention, how data is shared, how data is secured, user rights to data, user control of data, etc]
- [Be prepared to show an example]

**First Iterative Design [8 min]**

Now I'd like you to act as a software developer for this feature on your product, Zoom. Taking the original Zoom Attention feature, let's say you were provided the following request to update and change the feature:

[Choose one of the following 3 based on the random order chosen for the protocol for this participant - can paste into a separate doc and share so the participant can refer to these]

- (1) You hear this feedback sample from users:
  - "We wish Zoom would display a notification or let people on the call see whether it's enabled."
  - "People could be just doing other things and not have the meeting window up. It doesn't mean they're not listening."
  - "Attendees should know how accountability and performance are potentially being 'graded' along with the limitations of them."
  - "So, my 'I'm actually watching this on a different screen' won't actually fly for much longer?"
- (2) You hear this feedback sample from users:
  - "That's kind of intrusive in a way that often doesn't happen in the physical space or in the physical world. Employers have always had incredible control over how employees spend their time, but the technology makes it faster, more invisible and more sophisticated."
  - "If you have to track people to make sure they pay attention during the meeting, the meeting is pointless and too long. Meetings that are short and packed with useful info nobody wants to miss, are well-attended."
  - "The 'attention-tracking' feature [...] was appropriate in some business contexts, but for many new consumers, it presented a privacy conflict."
- (3) You hear this feedback sample from users:
  - "Everyone in the meeting room, the host, the speaker, and the attendees should be able to see who's paying attention."
  - "I think the issue is not that Zoom knows if its application window has focus, but that it \*reports\* focus state to anyone other than the user."
  - "Android users could look at their notifications but not iOS users [...] if you were on Android, you could bring down the Notifications tray for an unlimited amount of time since it didn't fully cover the Zoom app like the iPhone's one did."
  - "The feature creates a scenario whereby you could be penalized by an employer for doing job-related things, such as checking your notes or updating a memo during an important meeting."

Now I would like you to draw up a change to this feature that would address this request (and any other needs that you might think of). If you feel your previous idea is sufficient, please explain how it addresses this request.

- [Based on the ideas, ask for follow-up explanation, such as with an info flow, user storyboard, privacy storyboard, system architecture diagram, etc.]
- [Also consider probing about what data is collected, data retention, how data is shared, how data is secured, user rights to data, user control of data, etc]

How confident do you feel this product will satisfy user needs and concerns (on a scale of 1 being not at all confident to 5 being extremely confident)?

### **Follow-up Iterative Designs [8x2 min] (Repeat x2)**

Now let's say you as the developer also receive this feature request:

[Choose one of the 3 listed above based on the random order from the protocol for this participant - can paste into a separate doc and share so the participant can refer to these]

As before, please draw up any fixes or feature improvements you would add to address this request (and any other needs you might

think of), on top of your previous ideas. If you feel your previous design is sufficient, please explain how it addresses this request.

How confident do you feel this product will satisfy user needs and concerns (on a scale of 1 being not at all confident to 5 being extremely confident)?

### **Closing [1 min]**

That's all the questions I had for you today! Is there anything else about your ideas or thoughts about anything we've talked about that you didn't get to mention yet?

Thank you very much for your time! I'll collect your information for paying out your incentive. If you would also like to stay in touch and see the potential results of this study, let me know and I'll reach out once I have updates to share!

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009