# Title:  Strict Inclusion Closed Group Reverse Lookup (SICGRL) Attack

# Introduction

We have found a vulnerability in Facebook's Group system. This vulnerability could be exploited in a manner that results in the loss of life. It should be evaluated at the highest priority by Facebook's security and privacy team.

Facebook promotes it's Group functionality as a platform to create a safe place for vulnerable populations[1] to seek support from peers. Despite this, any Closed Support Group on Facebook reveals a list of members to all users. For any Facebook Group with strict inclusion requirements, this functionality amounts to publishing a personal fact about the user, which is nonpublic user information..

This document demonstrates this dangerous vulnerability, which will be referred to as 'Strict Inclusion Criteria Group Reverse Lookup Attack' (SICGRL). This is a system-wide vulnerability with a range of components:

- Third parties can easily scrape real names, locations and contact information of closed group members, even when its members did not explicitly consent to join the group.
- Group membership does not always equate that all members share some specific feature, this is only true when the group requires an inclusion criteria (i.e. all members must demonstrate that they have HIV).
- It is possible to write a script that searches for all groups that have an inclusion requirement like this, making it possible to attack all groups that have this vulnerability in an automated fashion.

The following sections will explain the scale of impacts, reproduction steps, classes of attack, and critical recommendations to address the problem. We will also detail how users can be added to groups without their consent, effectively "outing" them for a suspected personal status or, even more sinisterly, as a mechanism of "accusation" of a particular vulnerable status.

We request that Facebook create a new type of Group that allows support groups to advertise their existence but not expose their membership list. We would like to introduce new permissions to individual Facebook users, so that they can control when their membership is exposed in Groups. These recommendations, as well as the more specific recommendations we make later in this document, are all designed to allow users to consent to having their information released from Facebook Groups that they join.

Unaddressed, this design flaw in Closed Groups may have dire consequences for vulnerable populations around the world that are using Facebook's group features. In worst-case

[1] First Communities Summit announcement, June 2017

scenarios, a scaled SICGRL Attack could lead to a breach of health information and other private information, which could lead to irreparable harm and even loss of life, at an unprecedented scale.[2]

# Description & Impact

There is a large subset of Closed Groups on Facebook that only accept members who demonstrate that they have a certain personal characteristic (i.e. “strict inclusion criteria”). The shared characteristic could be a diagnosis, a genetic marker, a gender identity, a sexual preference, a clinical diagnosis, etc.

Many Facebook Closed groups that adhere to strict criteria for inclusion are created for the purpose of support, and are intended to be a ‘safe space’ to share private information among peers.

If any Facebook user can download the membership list of a Closed group with strict inclusion criteria with this personally identifying information, it is equivalent to releasing that given personal characteristic to the public without consent.

Some of these support groups have been created specifically for patients who share private health information. Group administrators use use Group admin features to ask screening questions as users apply for Group membership. For vulnerable populations these forms typically ask for confirmation of clinical status before joining a group.

The tool Grouply.io offered a mechanism to automate the download of the membership list of a closed group. This tool was capable of downloading the real-name, location, employer information and email of the members of Closed Groups. This method, applied against a vulnerable population that has chosen a Closed group on Facebook could result in a life-threatening and critical data breach.

Grouply.io was taken down as of May 15, however the tool is still working for anyone who previously installed it on their web browser. There are also a range of other similar web scraping tools that can achieve the same purpose to extract the entire member list of these groups such as Dex.io, Scraper Chrome Extension and any number of other generic browser based scraping extensions. Further, the underlying vulnerability has not been remediated; re-scripting the tool would be a trivial effort.

[2] Breach Portal: Notice to the Secretary of HHS Breach of Unsecured Protected Health Information

# Reproduction Steps

This section details the steps that Facebook can take in order to demonstrate that this is a real problem. Unfortunately, it also represents a how-to guide that could be followed by a malicious actor in order to take advantage of this flaw. This is the most important section to not discuss publicly until the problem is completely addressed.

There are two major steps to take advantage of this problem as scaled attack. First, there is the automated download of a single Closed Group's membership. This is, by itself, a vulnerability. But it only impacts a single group at a time.

However, it is possible to scale this vulnerability to every impacted group with a second step. The second step is to search across all Closed groups for groups that meet the following criteria:

- The groups require strict inclusion requirements for their members.
- The strict inclusion requirement represents some vulnerable status.

Facebook's native search functionality can be used at scale to find groups that are vulnerable programmatically. This changes the vulnerability from a "group at a time" to a method that could specifically target all impacted groups, and scale in a matter of hours in a data leak that would impact tens of thousands of groups and millions of users.

Both of these steps do not require special software, and can be implemented inside any modern browser that allows for substantial custom javascript to be used (which is true of most browsers).

## 1. Steps to scrape the group membership for a single vulnerable Closed Group:

1. Log into Facebook
2. Search for any Closed Group with strict inclusion criteria. Examples include:
   - REDACTED: Link to Group for abused women
   - REDACTED: Link to a group supporting diabetic patients with sexual dysfunction
   - REDACTED: Link to a group for women choosing to get Mastectomies.
   - REDACTED: Link to a group for people impacted by addictions
   - REDACTED: Link to a group for HIV support

    - REDACTED: Link to a group for supporting a personality disorder.
    - REDACTED: A group supporting Huntington's Disease.
    - REDACTED: A group supporting Genetic Conditions
3. Click the “Members” link on the top left of the page, just under the “About” link
4. Scroll down and record the members’ real names for the Group’s entire list using any of the following methods:
    - Data mining tool such as Grouply.io
    - Custom javascript
    - Excel spreadsheet or
    - Screenshot
    - Pen and paper

## 2. Automated steps to programmatically identify vulnerable groups and harvest membership lists from Facebook’s User Interface:

1. Create a list of clinical terms, METHOD REDACTED
2. Login to Facebook, use a javascript tool to repeatedly type the “target criteria list” terms into the search box.
3. Visit each group that returned from a “target criteria list” term
4. Click the link to “join” the group.
5. The contents of a typical popup look like this from the example **group**: REDACTED

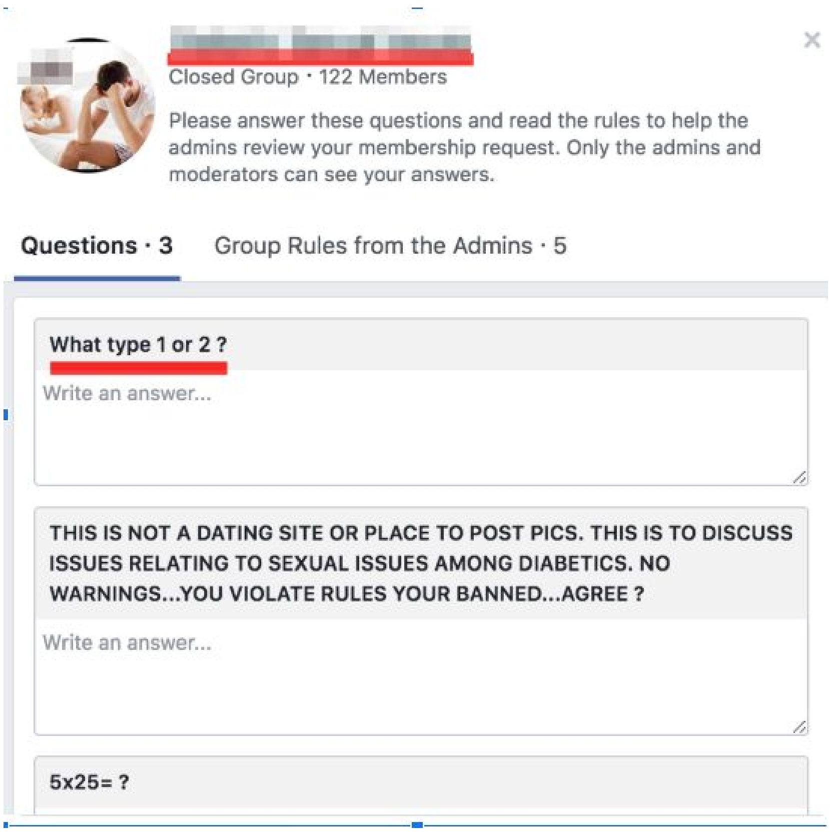

[redacted]

Closed Group · 122 Members

Please answer these questions and read the rules to help the admins review your membership request. Only the admins and moderators can see your answers.

What type 1 or 2 ?

Write an answer...

THIS IS NOT A DATING SITE OR PLACE TO POST PICS. THIS IS TO DISCUSS ISSUES RELATING TO SEXUAL ISSUES AMONG DIABETICS. NO WARNINGS...YOU VIOLATE RULES YOUR BANNED...AGREE ?

Write an answer...

5x25= ?

6. Programmatically or via semi-automated method such as Mechanical Turk search for indications that the group excludes people who do not have the clinical condition. Examples of this include term and phrases like:
   - “Group only for”
   - “Required in order to join”
   - “Must have”
   - “Users must be”

   etc

7. Using the “screening questions” given to new members is easily scrapable via a chrome extension, or by javascript commands entered into the console.

Using any of these approaches it would be relatively simple to generate a list of millions or tens of millions of Facebook users, along with a corresponding list of verified clinical or personal characteristics.

## 3. Moving support groups from “Closed” to “Secret” is NOT an effective remediation

Currently, a Facebook support group with a strict inclusion requirement can only protect their membership lists by becoming a “Secret” group.

While this does ensure that membership is private, it is not an appropriate choice for support groups. The whole point of a support group is to allow vulnerable people to **find** help from others in their community. Secret groups cannot be viewed in Facebook search results. In fact, a “Secret Group that can be found” is precisely what we are suggesting that Facebook enable to fix this vulnerability.

Current Facebook support groups will be forced to move to Secret if Facebook does not intervene and fix this vulnerability. Some groups are already doing this. In reality, this response will create a problematic void in Facebook’s group landscape which could be used by nefarious actors to setup “false front” operations as “replacement” support groups. One specific Facebook group that Facebook continues to recommend as a “safe” place is actually a marketing effort by a commercial Rehab facility. Many Facebook Groups for patients and vulnerable populations are already fronts for for-profit organizations that are marketing themselves.

If legitimate community-led Facebook support groups are forced to remove themselves from search, in order to protect their members privacy, they are likely to be replaced by either:

- Other well-meaning community members who have no idea that by “adding” community members to their Closed Groups they are “outing” their status as a vulnerable person.
- Or false-front corporations who do not care that they are deceiving their users by representing the group as a “private” community. These users may even use their membership in the newly Secret groups to add members to the Closed group so that they can “plausibly deny” that they were marketing using the Secret Membership lists.

In short, the only currently available remediation option available to Facebook Group Administrators (i.e. switching to Secret) will only serve to further violate the privacy of Facebook users at scale.

# Threat Modeling for SICGRL Attack

The potential harm to members of the Closed groups can take a variety of forms: a person could lose a job, or could lose insurance, or could lose a friend or a spouse as a result of the breach. Some percentage of affected users have committed suicide following similar data breaches.[3] There are two recent data breaches that may serve as precedents as we model potential threats to Facebook Users.

## 1. Ashley Madison Breach

If this vulnerability were to be made public before it is fixed, we expect that we would see malicious activities similar in kind to what happened following the Ashley Madison data breach.

- In that case, the malicious actors who were willing to attack individuals did not know how to hack Ashley Madison directly, but once the data appeared on the dark web, they were willing to use the data to target specific individuals, with blackmail threats and/or simple outings.
- Currently, if a malicious person wanted to harm a specific individual who was a Facebook user, it is extremely unlikely that they would attempt to "Download the membership of every Closed Facebook Group" in order to figure out if their intended target was a member of any of them.
- However, with the Ashley Madison breach (the only precedent of which we are familiar) one malicious party downloaded the data and put it on the dark web, a second malicious party made that data searchable, and still further malicious parties were then able to simply type the name of their targets to see if the Ashley Madison site had "dirt" that they could use to hurt that person.

## 2. Grindr Data Leak of HIV Status

The recent Grindr data leak may also serve as an precedent.[4] Grindr was considered to be a 'safe space' for gay men to convene., However, due to real-name data being released to third parties, in combination with a HIV status, Grindr serves as an example of harm when trust is lost.[5]

[3] Pastor Outed on Ashley Madison Commits Suicide, September 2015
[4] Grindr is revealing its users' HIV status to third-party companies, April 2018
[5] Grindr was a safe space for gay men. Its HIV status leak betrayed us, April 2018

- Men felt comfortable posting HIV status on Grindr because they perceived the site to be for gay men only. Therefore, they did not expect information to be accessible to third parties, even for the purposes of advertising or technology monitoring.
- For patients and other vulnerable populations that organize large support groups, there is a similar perception of ‘safe spaces’ via Closed Support Groups on Facebook.
- A breach of this data might cause vulnerable people in need to avoid Facebook Support Groups that might provide life-saving information.
- Similarly, loss of trust in Facebook Groups as a safe space to convene would represent the loss of a vital resource to support peers when they frequently do not have support through the healthcare system.[6]

## Estimations of Impact for Closed Facebook Groups

Based on our analysis of a small sample of Closed Support Groups data, we estimate that there are between five thousand to fifteen thousand English-language Closed groups that require strict inclusion criteria( i.e. evidence of specific clinical or personal status before joining the group.)

We expect that there are a roughly equal number of these types of Groups in other languages. Any Closed Support Group with strict inclusion criteria can range from hundreds to thousands of members.

**We may quantify the risk of harm as follows:**

> Multiply (1) a given Facebook Closed Support Group’s percentage chance of having each individual targeted for harm, by
> (2) the likelihood that they would be targeted, by
> (3) the likelihood that an attack would be successful.
> This would equal the number of times that such attacks might succeed.

These numbers may vary wildly from group to group. Facebook has support groups that number in the hundreds of thousands, and can be as small as two or three people. Many of these support groups are unlikely to be targeted for malicious purposes, while others are frequently the targets of physical violence. Some groups are so vulnerable that it would be very easy to harm them, while others would be very difficult to attack.

For example, REDACTED groups that share information about REDACTED as part of inclusion criteria for the group can be the target of predatory marketing, health insurance denial, discriminatory practices, or blackmail. These threats are realistic, we know that Ashley Madison

[6] Peer-to-Peer Health Care. Susannah Fox. February 2011.

users were blackmailed at scale, and we know that “patient shoppers” for rehabilitation services are already marketing services to Facebook users. It is likely that these marketers are already taking advantage of this vulnerability.

By our estimates, the chances of harm (defined below) are very low per user: in the range of 1% or to 0.1% probability of harm in the event of a large-scale breach. Given the population size of these user groups, the chances that users are harmed can be illustrated with this example:

- We have verified there are approximately thirty thousand people in the vulnerable Closed Groups with strict inclusion criteria for REDACTED.
- On the conservative side, let’s estimate 50% of the thirty thousand are part of a scraped list that discloses their REDACTED status via reverse lookup attack.
- 0.1% people successfully targeted = 15 people harmed

## Limitations of Threat Modeling

We recognize threat modeling has limitations. Threat, harm, scale, and impact are all difficult to forecast with precision.  There is far more deviousness in the world than can be readily imagined. It is entirely possible that we have missed something.

# Classes of Attack:  In Order of Severity

Based on the threat modeling we outlined in the previous section, there are multiple ways that a reverse lookup attack can harm Facebook users who are part of Closed support groups.

This section outlines several classes of potential harm, organized by level of severity, as we expect this kind of “breach escalation” to be a component of almost all classes of attack. This makes the classes of attack that we list below only the “last stage” of the attack cycle, representing “how it would be felt” by the targeted Facebook users.

## Class 1: Potential physical harm and loss of life

**This vulnerability can lead to loss of life.**  Easy access to vulnerable populations’ contact and location data by hate groups, blackmailers, and malicious hackers could cause physical harm to specific individuals:

- For REDACTED support groups such as REDACTED (2000+ members), a reverse lookup attack described in this report can easily expose the entire membership of this group, including location and contact information. These could allow targeting by anti-Semitic hate groups. (REDACTED mutations in general, and a short list of specific mutations particularly, are significantly more common among people of Jewish heritage.)
- For Closed groups who are HIV positive, blackmailers might threaten individuals by extorting money to keep a member’s HIV status from becoming public (For example REDACTED with 200,000+ members)
- On a global scale, the reverse lookup attack can easily be exploited to make lists that include contact information and locations of ethnic or religious minorities and/or political dissidents who are participating in Closed groups with strict inclusion criteria.
    - This is most dangerous, where all members of a Closed group can be targeted.
    - It is possible, for instance, for all members of a Closed group of gay men in REDACTED (REDACTED) to be outed and then face state-sanctioned torture and/or murder as a result of being a member of the Closed group with that inclusion requirement.
    - Because of language and cultural barriers we do not yet have an accurate way of knowing how many such extremely vulnerable groups are on Facebook.
    - A worst-case scenario would be the first mass casualty event due to a data leak.

## Class 2:  Exposure to financial discrimination practices by insurance companies, employers, and/or credit agencies

- Insurance companies could use scraped data from Closed Support Groups with strict inclusion criteria (such as REDACTED status) for underwriting purposes without individual users' knowledge or consent.
- Employers and recruitment agencies could use scraped data from Closed groups with strict inclusion criteria to make decisions about job candidates.
- Credit agencies could use the health status of people participating in Closed groups to make decisions about loans.

## Class 3: Exploitation of Closed support group data by organizations offering spurious treatments

- Scraping these Closed Support Groups could create large-scale lists of patient registries for third parties. There is a great deal of economic value in some of these lists with strict inclusion criteria, and these lists can be used to target, market, and spam vulnerable patient communities.
- We also believe that recent news about abusive rehabilitation centers' recruiting practices could already be leveraging this data leak.[7]
- Scraped data could be used by drug companies and medical device companies to market treatments -- legitimate or otherwise -- directly to patients.

While some forms of marketing in this manner may be illegal in the US and the EU without a patient opt-in, the fact that patient identities are made available to unscrupulous marketers due to this vulnerability must be addressed.

## Class 4: Leaking of data to random companies, with random results

These companies have no specific agenda to harm vulnerable populations, but they also have no legal obligation or the practical understanding required to ensure that data is not further shared downstream. Class 4 is essentially "Classes 1, 2 and 3 by proxy".  This opens the door for further misuse of these sensitive data. Grindr's recent communication of HIV status of its users to third-party companies, is an example of a company inadvertently leaking data without fully understanding the potential consequences.

[7] This type of marketing is already a known problem on Facebook, with the recent release of Predatory Behavior Runs Rampant In Facebook's Addiction Support Groups, by Cat Ferguson at The Verge.

# This a system-wide design flaw

The potential for a reverse lookup attack isn't merely a technical vulnerability. Rather, it is a system-wide design flaw in the platform's Group functionality. In the 2011 FTC settlement, Facebook commits to "*not misrepresent in any manner, expressly or by implication, the extent to which it maintains the privacy or security of covered information*". This 2011 settlement goes on to specifically cover the interactions between users and "Third Parties," which are directly pertinent to this design flaw.

There are a range of issues that run contrary to the principle of privacy by design mentioned in the FTC settlement, which we'll explain in this section:

- Facebook's real-name policy represents a breach of 'Privacy by Design' when combined with certain Closed Group features.
- The choices given for privacy settings create a dilemma for Administrators that funnels vulnerable populations into the Closed Group option *by design.*
- Closed Group features on Facebook enable misrepresentation of privacy to Administrators, and to members who join for support.
- Force-adding group members can cause users to be "outed" by membership without ever having seen information about the security and privacy settings of Closed Groups.

## The Design Flaw of Closed Group Settings

Facebook allows Closed groups to be explicitly listed as "Support" groups, and has encouraged patient communities on its platform for years. Led to believe they could create a safe space to offer support and help for their fellow patients, Administrators of Closed Support groups for patients have assured members of these groups that their data is "private" and "secure".

Indeed, the word "Closed", which Facebook chose to summarize the privacy group functionality of the group, gives an assurance of "privacy". There is nothing in the word "Closed" that gives an indication that Membership would be shown to the public, or that third parties can easily generate lists of the group membership with each member's contact information. While Facebook's information is clearly displayed to users on *some* screens regarding Closed memberships, subsequent screens, controlled by Facebook group administrators might give **exactly the opposite indication**, promising users that the Group will **respect their privacy and keep their secrets**. While not every Closed support group has made this mistake, they do frequently communicate in a way that conflicts with what Facebook says. Some groups do this even while requiring a vetting process that ensures that users are exposed.

# Real-Name Policy

Facebook requires users to use their real names on the platform. When combining this policy with certain features outlined below, users real names can be associated with certain group identities that reveal clinical information (HIV status, REDACTED Status, Cancer diagnosis) without consent. Thus, when added to a Closed group with an inclusion requirement, members are outed with their real-world identity without consent.

# The Group Administrator's Dilemma

Closed Support Group Administrators must choose between the only options available to them. Group Administrators want to facilitate conversations that cannot be public. When creating a group, new administrators must make a choice between three options:[8]

- Public
- Closed
- Secret

**"Public" groups are not an option for support group Administrators.** For example many REDACTED cancer support groups such as the REDACTED LINK include photos of Mastectomies and other incredibly personal information (in the specific case of REDACTED). Photos and deeply personal information shared among peers in this group cannot be public information.

**"Secret" groups are also not a viable option.** Specifically, Secret groups do not show up in Facebook search results, because new people in need of support cannot find the group to request membership. Therefore 'Secret' group settings are also not a viable option for these patient support groups. Moving groups to secret means that new members have no mechanism to find the groups available to them and request to join.

**The only option left is a Closed group, which leads us to this current vulnerability.** As a result, Facebook funnels patient support community to the "Closed" group type as the only option between Public and Secret. There is clearly accessible functionality to ensure that membership status in Public groups is not visible on a given account's profile page. However, there is no equivalent setting to prevent the reverse lookup of membership in Closed Groups.

# The privacy risks are not apparent to Administrators

We do not believe that Facebook Group Administrators are maliciously or deliberately putting their users into a position where their information can be leaked. Rather, many of these

[8] What are the privacy settings for groups?

administrators are patient advocates who have worked for years to cultivate a supportive and safe space for members of their support groups without being fully informed of the privacy risks and implications. While the group administrators who create Facebook Groups are made aware of the privacy options, subsequent Group Admins may have never been informed about the options of Group type and the implications thereof.

Once a Facebook Group administrator has made a decision to pair a Closed group with an inclusion process, there is no way for an individual Facebook user to roll back that decision in their personal settings. The combined effect of being able to tweak visibility settings on Public groups, and not for Closed groups, encourages a user to assume that a Closed group membership status is "private", when the opposite is true.

Facebook has delegated substantial portions of the burden of communicating and enforcing privacy standards to support group administrators, who have no training in privacy or data protection standards, no understanding of risks for this reverse lookup attack, and no technical understanding of how data from groups has been easily scrapable by third parties. Further, the group administrator User Experience provides no specific clues to the privacy related decisions that Group Administrators might make on an ongoing basis. The interface does not adequately inform people who have no other source of context of the ongoing implications of the designs of the underlying privacy infrastructure.

So far, every Group Administrator that has learned about this issue has been shocked to find out that membership data is readily available to "anyone".

## Force-Adding Members to Closed Groups without Consent

More concerningly, Facebook users can be added to groups by group members without the permission of the individual user. We will refer to this group as the "force added" exposed user group. These Facebook users did not consent, and may not even be aware that they are members of Closed groups.

These users can be "outed" by membership without ever having seen information about the security and privacy settings of Closed Groups at all.

For instance, a user clicks "join group" but decides, based on learning that membership in the group can be seen by "anyone," not to join the group, and further decides not to fill in the form that requires "inclusion proof." Then the user leaves the form, and wanders off to some other part of Facebook. But later, the admin saw that the user in question started joining the group, but did not answer the survey, but decides to add that user in any case. This "force add" feature can, in many cases, publicly "out" the added user for a personal or clinical issue.

In short, users have been denied the ability to limit Closed group memberships to only those groups they affirmatively consent to join, the underlying design of the groups is contrary to “privacy by design” and users have been both implicitly and explicitly misled about privacy settings.  All of these problems are a deviation from both the letter and the spirit of the FTC settlement.

## Privacy Design for Closed Groups

If any Facebook user, even those outside the Closed Group, clicks on the membership list of any Closed Group, that user has the ability to view, mine, and scrape the complete membership of that group without consent from the individual user or the group administrator. This is true even for Closed Groups that are explicitly labeled as support groups.

Facebook users who join Closed groups frequently viewed two different messages at the time they joined. One message came from Facebook that indicated that their membership was going to be public information, and the other message came from Group Admins who indicated (incorrectly) that their information would be kept private. These mixed messages are a result of the limited choices between Public, Closed, and Secret groups that are explained in the ‘Group Administrator’s Dilemma’ section above.

Either way, the result is a group of users who have had their clinical or personal information made public by consenting to a “mixed message” from Facebook’s system-as-a-whole.

# The FTC settlement is not just a contract, it is a software specification. [9]

These issues in the design of Facebook’s privacy settings on Closed groups represent a vulnerability not in the source code of Facebook's servers, but in the system that Facebook agreed to implement with the FTC, and by proxy with the public and its users.

We urge Facebook to take this issue seriously and take responsibility for the fact that (a) there was no reasonable appropriate option available to the relevant Facebook group administrators and (b) very few end users participating in these groups intuitively understand the implications of the reverse lookup attack.

Essentially, while the Facebook software stack might technically be operating as advertised, the Facebook system, which is what the users experience, is profoundly broken and vulnerable, making substantial corrective steps reasonable and necessary.

[9] Facebook Settles FTC Charges That It Deceived Consumers By Failing To Keep Privacy Promises

# Critical Recommended Actions:

## 1. Fix the current issue by mandating a secure option

We recommend that Facebook change the security settings available to Groups in three ways. Once those improved settings are available, then users' settings should automatically changed to the more private mode. This way, Facebook users could later choose to consent to more relaxed privacy settings, if this is what they want.

The three configuration changes are:

1. Create a new "Private" Group Setting that is searchable in the Facebook search, but does not make the membership list public (other than the count of users). This option would provide sufficient privacy protections for most Closed groups with strict inclusion criteria that focus on sensitive topics and/or support groups.
2. Create a new configuration panel for individual Facebook accounts that allows users to determine which Closed groups can display their information for reverse lookup. This choice should be honored for reverse lookups on the Groups pages (which should be modified to merely list the number of hidden users) and should also disable the ability for "searching closed groups of which my friends are members" feature of the Facebook search.
3. Do not allow any user to be "force added" to any Group without their consent.

After the new group type is available and functioning properly, we recommend that Facebook migrate every currently Closed group into the "Private Group" status. Group administrators should be informed about this vulnerability, and why this change was made, and then be given six months (or some similarly reasonable time) during which they can change their Group type back to Closed, if they decide that publishing their membership list is in the Group's interest.

Similarly, the individual settings for "allowing reverse lookup" for all Closed and Private groups should be set to "no" for every member of either a Closed (or the new Private) Groups. This means that even if a Group Administrator chooses to make their group "Closed" again, it would require user intervention to allow their identity to be listed in a reverse lookup.

It might not be necessary to make this change for every Closed Group; the issue arises only for groups that have populations that are vulnerable for some reason. However, limiting the change to certain groups would require that Facebook be in a position to accurately determine which Closed groups have established strict "inclusion requirements." This could be inordinately difficult, given that many groups choose to include based on content of messages, etc, rather than by using the sign-up forms. For this reason, we recommend that Facebook make this change automatically to all Closed Groups, and allow group administrators to affirmatively elect that they really want to have their membership lists published.

It might be possible to limit the automatic move from Closed Group to the new “Private” Group setting to those groups that are specifically labeled as “Support Groups” in their Group Type selection, or Groups that have no label at all. Alternatively, it might be effective to only auto-migrate groups that have ambiguous types. For instance, a “Family” or “Parenthood” group could easily be chosen by parents of rare-disease children, but it might be safe to assume that “Video Game” groups could safely remain as “Closed”. There could be other methods used to determine if a given group needs to move to Private or can safely remain Closed. These might include the use of tags, the use of highly specific questions (indicating an inclusion requirement) or other metadata about a given Closed Group. Using any method other than simply migrating all Closed Groups to the new Private setting should only be considered with extreme caution, since a false negative on the decision to move could be disastrous, while a false positive (making a Private a group that really should stay in the “Closed” status) is merely a frustration.

This is the only way that we can see to correct the problem for all impacted parties without also informing malicious users that specific Groups may be vulnerable.

## 2. Automatically set the privacy of “display messages”.

In a similar manner, when a group moves from any other group type to public, messages made by an individual should be hidden, in accordance with the group setting when they signed up with the group.

Other data privacy settings related to Groups should automatically follow this “presumed preference for private” as the status of Groups are that are reset to different types of privacy levels.

## 3. Inform users about already-leaked data.

The current data leak mechanism requires a user to scroll through thousands of individual names in a given group. While typical users might have a passing interest in knowing which of their friends are in a given group, they are unlikely to scroll to the bottom of a page with 1000+ entries, especially when each scroll requires a slow and distinct AJAX call.

Even more rare is the case where a single account would legitimately subject themselves to this scrolling experience in order to reach the end of multiple lists that contain thousands of members. A legitimate version of this type of behavior must be rare. The fact that the legitimate version of this behavior is very rare creates an opportunity for Facebook to inform users about already-leaked data.

Facebook should immediately review its logs to determine which users were data mining this information. Especially if users were doing this in a manner that indicates that they did not understand that doing this was a violation of Facebook terms of service (i.e., they were not using fake accounts, they were not changing their IP address, etc.). Given that other social networks encourage this type of data mining, it is not unreasonable to assume that good actors could simply be confused about what was allowed. In those cases, Facebook should reach out to the scraping parties and request that they delete the data that they have.

Facebook should also estimate, where possible, when more sophisticated black-hat actors have sought this data at scale. Facebook uses algorithms to determine when it is being “scraped” and to model which adversaries are acquiring what information. Facebook should release a report estimating how many malicious parties potentially have access to this data. Moreover, it should calculate specific details on a per-facebook account basis and inform specific users how often their data may have leaked using this mechanism.

If at all possible, Facebook should seek to inform when its users have been targeted by at-scale trawling of closed membership data.

## 4. Ensure that membership status of any user will not be published in a manner contrary to user consent.

Members of Closed Support Groups on Facebook need new privacy settings to choose whether their group membership is publicly visible or private. The current privacy settings available to group members are not adequate. Currently, groups with less than 5000 people can choose to become more public, by choosing to change a group type from Secret to Closed or Public, or from Closed to Public.

- Individual group members need the ability to easily determine whether their membership in that group is public information.
- But more importantly, if a group administrator changes a group from a type that does not reveal membership information (i.e. Secret or the new Private) to a group that does display membership options (Public or Closed), then their individual membership display option should be automatically set to “not shown,” to match the setting when they established the group.

## 5. Report this problem to assessors and employees responsible for privacy, as well as appropriate regulators.

The 2011 FTC privacy settlement required that Facebook submit to a qualified, objective, independent third-party professional, who uses procedures and standards generally accepted in the profession for the purposes of a privacy and data protection review. This vulnerability and Facebook’s response are relevant to that assessment. We expect these materials to be provided to that assessor(s).

The same settlement requires that Facebook ensure “the designation of an employee or employees to coordinate and be responsible for the privacy program”. We expect that this employee will be similarly provided with this document and involved in the discussion about this issue.

Since 2017, Facebook has specifically recommended that its users contribute “stories” to support groups for the purpose of improving their health. For example, here is a video of Matthew Mendoza’s group, REDACTED. We cannot help but note that the name of the people posting to the LINK REDACTED in this video where blurred to protect their privacy, a dignity that the Facebook platform itself currently denies those members. This is embedded in a larger announcement about Facebook describing how groups are a safe place for vulnerable populations to seek support for health related issues:[10]

> *Matthew Mendoza, who started* REDACTED *Support Group. The group is a safe space for people who are experiencing or recovering from drug and alcohol addiction, as well as their friends and family, to offer support and share stories.*

As a result, we believe that Facebook qualifies as a Personal Health Record (PHR) under the FTC’s definition which reads:

[10] Our First Communities Summit and New Tools For Group Admins, Facebook Press Release, June 2017

> *A **personal health record** is defined as an electronic record of "identifiable health information on an individual that can be drawn from multiple sources and that is managed, shared, and controlled by or primarily for the individual."* [11][12]

Facebook has recommended specific patient support groups and, in general, those support groups have recommended that patients upload their health information to the group in the form of posts, so that other users can comment and help them with their healthcare issues. We believe that by explicitly recommending Closed Groups as "support" groups, in order to address healthcare concerns, that Facebook qualifies as a PHR.

To qualify as a PHR, there are two criteria for the FTC health breach notification rule that apply. These two criteria met by Facebook as follows:

- The first criteria to qualify as a PHR is that Facebook offers Groups functionality to support groups for clinical topics.
- The second criteria is for the information leak to qualify as a "breach" under the rule. The first specific criteria for a "breach" under the FTC rule is "acquisition of such (healthcare) information without the authorization of the individual".

We have provided evidence and examples in this document to meet the second criteria above. Closed Support Groups on Facebook with a clinically defined strict inclusion requirement qualify under the FTC health breach notification rule for a range of reasons that include the following:

- Grouply.io and similar tools can and have (at least once) been used to download the membership list of a Closed Support Group. These group members all share a common clinical fact about them.
- There are members of Support Groups who may never have explicitly consented to share the healthcare fact (i.e. non-public user information) in a public way.
- Given that other users can force-add members to Facebook Groups, there are at least some users on the platform for whom the download of their membership in a Closed Group with a clinically defined strict inclusion requirement represents a breach under the FTC health breach notification rule.

# Recommended Further Actions

This bug cannot be regarded as fixed until the Critical Recommend Actions, or their equivalents, have been completed. Once that has happened, these further actions can ensure that Facebook is making its Groups more sensitive to the privacy needs of vulnerable populations.

[11] Complying with the FTC's Health Breach Notification Rule

[12] Underlines added for emphasis, and are referencing emphasis of reasons why Facebook is covered by these rules from the FTC

## 1. Provide tools for group admins to handle privacy fallout

Currently, Facebook allows public, closed and secret groups to choose the options for a **Support Group.** This further encourages group administrators to choose Public and Closed Group types , despite inappropriate privacy settings

Facebook needs to offer settings that better differentiate between **Private Support Group** and **Public Support Group.** Improved privacy settings for support groups are necessary to enable users to make an informed decision about whether their nonpublic user information can be revealed to the public. The current group type selector looks like this:

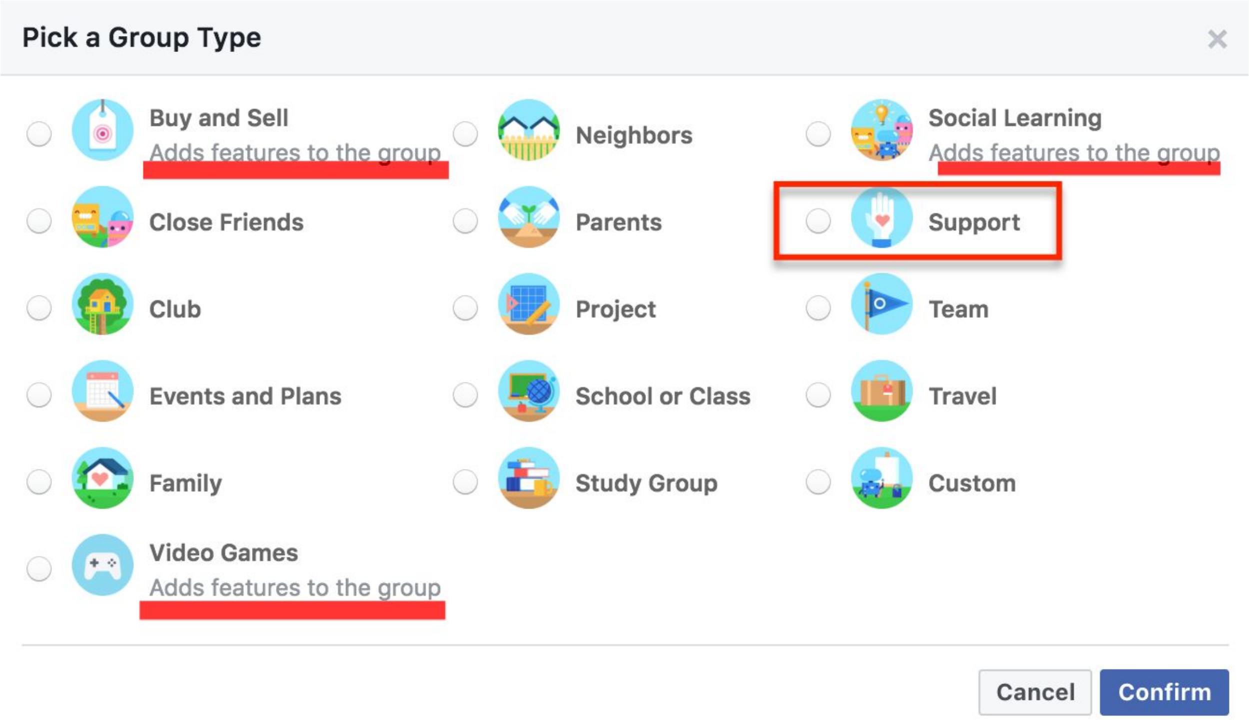

There is only one "support" type. There should be two -- "private support group" and "public support group" -- and the options to choose a "private support group" should only be available in "Secret" and our suggested "Private" group types. This will help make it clear that Public and Closed group types are not appropriate for a Support group that advertises itself as protecting the privacy of its members.

There are additional features available for groups categories including "Buy and Sell," "Video Games" and "Social Learning." There should be additional features available, especially to group administrators, in any "Support" group type. At a minimum, in order to deal with the potential fallout of Closed group memberships being public information, there should be a "priority" queue for reporting accounts that exhibit suicidal behavior, and for connecting vulnerable populations with the tools to access emergency international and national police resources, and potentially human rights groups that are capable of extracting people from dangerous environments quickly.

These options need to be in place so that if vulnerable populations have their group membership leaked, and there is a malicious attack on the members of a given group (or the members who reside in a specific location) there will be options for group members and group administrators to quickly coordinate help via a designated point of contact at Facebook.

## 2. Create an amnesty program for those who have previously scraped data

Facebook should offer an amnesty program for data scrapers who may have downloaded group membership data without understanding that this was a violation of Facebook terms of service. Parties would be eligible for this amnesty if they come forward with their data scraping activities, and:

- Are willing to identify which user accounts they created to scrape the data
- Are willing to detail what specific data that they had downloaded
- Are willing to detail any third party that they shared this data with
- Are willing to delete all downloaded group membership data
- Are willing to identify which patient support groups had their entire content histories scraped via the group API. (NOTE: Before the API was shut down in April 2018, the content of the entire group could be scraped by *any member* without consent of individual users or administrators.)

This amnesty would mean that Facebook would not pursue civil or criminal charges against these users for their actions. This offer of amnesty should last for 6 months.

Once the amnesty period is over, Facebook's security team should identify, and Facebook's legal team should pursue, all persons and organizations who violated the facebook terms of service at scale to violate users' privacy by exploiting this bug.

This "amnesty then accountability" mechanism should encourage merely ignorant data miners to self-identify and delete data they have acquired. They will also enable the Facebook security team to learn to pattern-match for when this kind of scraping is happening at scale. This will make subsequent and ongoing enforcement efforts more effective.

It is entirely reasonable that this kind of "data leak remediation" work should happen some weeks or months after the ongoing data leak problem has been corrected.

## 3. Provide resources to patient groups and vulnerable populations who are blackmailed, spammed, or attacked

In the event that this leak does result in attacks, there should be a mechanism for users to report such attacks, and to attain legal help from Facebook-hired attorneys when this is required.

We can provide evidence that patient groups are receiving unsolicited targeted marketing emails from third parties, based on their membership in a given Facebook group. Reporting unwanted marketing is a good mechanism for Facebook to leverage a reporting mechanism that will make “group-based” marketing enabled by scraping more difficult in the future.

Facebook should cooperate with, and in fact seek out the assistance of, law enforcement to track down those who experience assault, blackmail or other illegal acts as the result of this data becoming public. As long as Facebook users can demonstrate that (a) they were part of a Closed group that had a data leak and (b) they were otherwise cautious in their Facebook posts or (c) do not know their attacker personally (i.e. suggesting that they were “found” using a reverse lookup attack) they should qualify for this type of assistance.

Rather than seeking to ensure Facebook’s resources are “only spent on harassment that is Facebook-born,” Facebook should seek to provide automated tools to combat real-world harassment of vulnerable people for any reason, and particularly where it may be tied to an exploit of the vulnerability described in this report. Facebook is in a position to automate access to appropriate resources, in a manner that would make this social problem independently better. Facebook’s nascent efforts to help coordinate access to mental resources for people that post to Facebook indicating suicidal ideation (which is a good start) should provide a template for this set of features. Specifically, there should be some mechanism for a person to report that they have been threatened or harassed, specifically as the result of participating in a support group.

# Adherence to Facebook’s Responsible Disclosure Policy

We have made every effort to adhere to Facebook’s “Responsible Disclosure Policy.” In keeping with Facebook’s policy with regard to Responsible Disclosure, we did not know we were “hacking” until we did, and once we did, we immediately stopped.

When the patient advocate who found this vulnerability brought concerns to a professional security researcher, this researcher initially assumed that a third party tool such as Grouply.io restricted access to Group Admins only, and that this tool would not work for non-admin users. We assumed that it was a power tool for admins and initially explored its functionality with the assumption that it would act differently for different types of users. When we tried the tool through a non-administrator account, we did so assuming that we would merely be confirming that the tool was an admin only-tool.

Once we confirmed that the tool worked for *any* Facebook user, against the REDACTED group, we stopped using the tool. The Group Administrators of the REDACTED were informed about this issue.

Facebook has a policy of listing researchers who contribute under its Responsible Disclosure policy. For the purposes of that credit on that list, this vulnerability was discovered by the REDACTED community advocate who chooses to keep her identity private for safety reasons. She will be referred to as Moana Mononoke (http://twitter.com/MoanaMononoke) until it is safe to share her identity. This vulnerability was initially verified by Fred Trotter (http://twitter.com/fredtrotter/).

It is critically important that Facebook see that this is the only important motivation that we have for reporting this bug. As a result, payment from the Facebook bug bounty program would be donated to a charity focused on patient privacy of our choosing. Our goal is to ensure we do not muddy the clarity that we seek to express regarding the seriousness of this issue.

# Schedule for Action

We understand that implementing these or equivalent solutions may take considerable time. However, we expect the reasonable amount of time for Facebook to come to a conclusion about its approach to this problem will be within several days.

As representatives of patient communities, our primary ethical responsibility is not to Facebook, but to the millions of Facebook users who are releasing information that makes them vulnerable to harm. We will give Facebook a short period of time to make a commitment to our remediation proposal or one with a comparable outcome.  We will also make every effort to be responsive to Facebook and help work towards remediation and viable solutions.

We are very intentionally not specifying deadlines in this document in order to balance our need to be reasonable with Facebook and our ethical obligations to our community.  One the one hand, we recognized that what we are requesting is a significant change in the design of Facebook's Groups system.  On the other hand, thousands of complaints from the patient community about these types of issues have been ignored completely by Facebook in the past.

If at some later time, Facebook fails to demonstrate that it remains committed to fixing this vulnerability in a responsible timeframe, then we will implement our contingency plan.

Assuming that Facebook is taking responsibility, then we expect the date we release this information to the public to be a joint decision, made to correspond with the deployment of the fixes, in a coordinated manner.

## Contingency Plan

If Facebook has not indicated that it will treat this issue with as much diligence and seriousness as it would a similarly dangerous "pure" software vulnerability, and treat a comprehensive remediation as a priority, then we will implement our contingency plan.

Our goal will be to ensure the maximum number of impacted populations are protected, while still accounting for the average Facebook ''user's right to know about this issue quickly.

If we regard that a contingency plan is necessary, our schedule for releasing information about this vulnerability to the public will be entirely dictated by the needs of that contingency process, and Facebook will not have input into those decisions.

It is not our intention to be combative with this position, but this is necessary given the highly vulnerable nature of the people who are impacted by this problem.

## Multiple journalists are informed about this issue and its seriousness

These journalists have agreed to a temporary embargo of the information in order to allow Facebook make a good faith effort to remediate the problem.

# Conclusion

Unaddressed, this vulnerability could result in the largest loss of human life in history as the result of a flaw in a digital system.

If this problem is not addressed in a comprehensive manner by Facebook itself, it will fall to a coalition of vulnerable populations who will be forced to try and fix this problem on a group-by-group basis before word gets out that these vulnerable populations can be programmatically targeted.  It is unlikely that this fix can be implemented in a comprehensive manner, across all affected groups.

If Facebook does not act swiftly to address this vulnerability, those groups, and eventually the public, will never forget that Facebook created this problem due to its desire to grow at all costs, advertised explicitly that the platform was an appropriate place to host support groups, and, when confronted with the problem that it created, abandoned its most vulnerable users to suffer at the hands of the worst parts of humanity.  That cannot be what Facebook has meant by "connecting people."

**Please take this issue seriously and do something about it.**

## Vulnerability Found & Report Authored by

Moana Mononoke
Community Data Organizer[13]

Fred Trotter
Data Journalist, member of Health Care Industry Cybersecurity Task Force
CareSet Systems

## Vulnerability Validated By

REDACTED

## Report Reviewed By

REDACTED list of experts

Update Feb 16 2019: This document is now publicly available at:

https://missingconsent.org/downloads/SicGRL_initial_report.pdf

[13] About Community Data Organizing