

Reporting Problems:

A Survey Analysis of the Administrator's Noticeboard/Incidents on the English Wikipedia

Caroline Sinderson, Sydney Poore and Patrick Earley, Wikimedia Foundation

Harassment is a serious problem within digital platforms and online communities. To better understand harassment inside of the Wikimedia community, the Wikimedia Foundation is devoting resources and research to analyze this problem. In 2015, the Wikimedia Support and Safety Team conducted an extensive harassment survey to analyze how harassment affects the Wikimedia community and projects. Over the past year (2017), the Anti-Harassment Tools Team has created several [small surveys](#), [interviews](#), [research insights](#) and requests to analyze [different forms](#) of [harassment](#), which include where harassment occurs and what harassment mitigation can look like.

In an effort to examine more deeply the current ways to address incidents of user misconduct on Wikimedia wikis, in 2017 the Foundation surveyed users about the English Wikipedia's [Administrator's Noticeboard/Incidents](#), (AN/I) which is a space designated for conflict mitigation and problem solving. Problems reported can range from identifying [spam](#), identifying [sockpuppet](#) accounts, undoing [vandalism](#), conflict arising from [bad edits](#), and mitigating and addressing different kinds of harassment. Although AN/I is a key part to the conflict mitigation process on English language Wikipedia, there has never been a survey focusing specifically on that topic.

The Wikimedia Foundation respects that AN/I is a community-led and developed process. While this survey is intended to understand the community sentiments around AN/I, including possible areas for improvement, this report will not lead to immediate or imposed changes to AN/I from the Foundation. Such change would need to be backed by consensus from the volunteer community on the English Wikipedia.

METHODOLOGY

The Support and Safety team (SuSa) and the Anti-Harassment Tools team (AHT) surveyed Wikimedia users from across the globe who contribute to the English Wikipedia in order to collect their feedback about AN/I. The focus was not just on the most "prolific" users or those who use AN/I the most, but on all different kinds of contributors who use AN/I in a variety of different ways. and how they feel about the AN/I process. In an attempt to get feedback from users with diverse interests and backgrounds, outreach was done to

There were 136 respondents in total. While this is a small sample, the survey was designed primarily to gauge feedback on uses of AN/I, and not intended to be a scientific survey.¹ This means that while the data can give us some directional clarity for this group, it is not scientific and thus results cannot be expected to reflect the views of the entire community.


¹ https://en.wikipedia.org/wiki/Survey_methodology

The survey was broken into three parts: use of AN/I, satisfaction with AN/I and experiences with AN/I. The goal was to gauge what AN/I does well, what it doesn't do well, and what needs to be improved. 136 people responded and filled out the survey. The survey has 23 questions, including six write-in answers, and two questions asking users to rate multiple features on a five-point scale. Answers from the six write in questions were sorted, grouped, and labeled to organize the feedback into cohesive takeaways. Some answers received multiple labels as they provided multiple kinds of insight.

The graphs below reflect the breakdown of survey participants by level of experience, location, and gender.


How long have you been a Wikimedia contributor?

136 responses




Which region do you usually edit from?

136 responses



Which gender do you identify with?

136 responses



OVERALL RESULTS:

Overall, the survey from both the qualitative and quantitative parts, provided suggestions and feedback for how to improve AN/I. The survey also helped us create a breakdown of the different kinds of users who use AN/I, along with the different ways in which they interact with and use the noticeboard—from weighing in on cases to reporting. The users were divided into groups to better understand user patterns and behaviors—like those who participate or weigh in on discussions, those that file reports, those listed in reports, and those that have been punished or admonished by the AN/I process. To further analyze the feedback, categories of users, such as more experienced editors who have been on Wikimedia projects 5+ years, male participants, female participants, were grouped by answers to specific questions relating to how well AN/I handled specific cases.

Both the quantitative and qualitative results highlight two major themes: a) the types of problems that AN/I is best suited to resolve and b) suggested improvements to the process for reporting and mitigating case. Users are suggesting some form of change to the AN/I process and how AN/I functions, from a both policy and product standpoint.

KEY FINDINGS (QUANTITATIVE):


What is reflected in the qualitative results are how AN/I, the structure of the board and who is allowed to weigh in, handles certain kinds of cases better than others.

Satisfaction with AN/I

There are mixed opinions among respondents about how well the general process works to resolve AN/I cases. There is a stronger dissatisfaction expressed about *the way that AN/I cases are handled* rather than *a negative view of the AN/I type process*.

On a scale of 1-5, how satisfied are you with the way reports are handled on AN/I?

participants (n) = 134




1.47% did not select an option [?]

Only 27% of respondents indicated they were satisfied with the way that AN/I cases are handled.

On a scale of 1-5, do you agree with the general process of how AN/I reports work? (e.g report, discussion, and then decision?)

participants (n) = 134



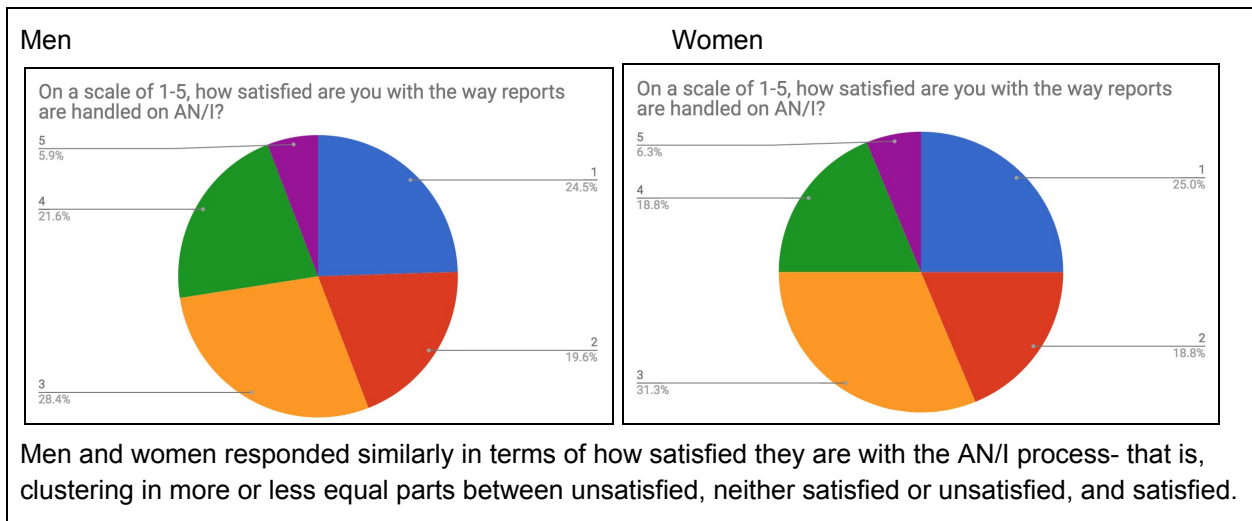
2.21% did not select an option [?]

More than 50% of users agree or strongly agree with the general process of how AN/I reports work.

There are demographic similarities amongst the three groups, "users who are dissatisfied", "users who are neutral", and "users who are satisfied", which also are the same as the overall demographics of the survey. The majority of users in all of these groups are male, experienced editors, and hailing from North America and Western Europe.

Around **12% of the participants were women**, which is roughly proportional to the [actual distribution](#) of male and female editors in the population of editors. The 16 participants is a relatively small sample, nevertheless, there are some findings worth surfacing, though further research is still needed.


Satisfaction with how reports are handled on AN/I:



Avoided reporting an incident because it would not be handled appropriately

Have you avoided reporting one or more incidents to AN/I in the last 12 months, because you did not think it would be handled appropriately there?

participants (n) = 134



- 1 – No (40.30%)
- 2 – Yes, one or two times (26.87%)
- 3 – Yes, between three and ten times (17.16%)
- 4 – Yes, more than ten times (8.96%)
- 5 – Other (free text) (6.72%)

2.21% did not select an option [?]

More than half of respondents said they have specifically avoided making a report on AN/I because they were afraid it would not be handled appropriately.

More than half of respondents said they have specifically avoided making a report on AN/I because they were afraid it would not be handled appropriately. Reason for interacting at AN/I

Self reported reasons to participating on AN/I

Self reported number and reason for interacting on AN/I




Type of Participation on AN/I by Experience Level of Editor:

The survey findings show that participants with 5+ years experience are weighing in on AN/I and reading AN/I more often than they are filing reports. 77% of these respondents are weighing in, to varying degrees, on cases they are uninvolved in. And 61% these respondents visit AN/I more than ten times a year to read or follow reports that do not involve them.


How Often More Experienced Editors Weigh In On AN/I

Count of How often have you been an uninvolved participant in discussions on AN/I in the last 12 months?



How Often More Experienced Editors Report Incidents to AN/I

Count of How often have you reported incidents to AN/I in the last 12 months?




While there were not many women responding to the survey, their likelihood of being involved in or reporting an incident in the past year was effectively equal (i.e. 69% for women and 71% for men).

Avoided participation at AN/I because of fear of retribution

Have you avoided reporting an incident or taking part in a discussion on AN/I in the last 12 months, because you were afraid of retributions of any kind?

participants (n) = 134



- 1 – No (56.72%)
- 2 – Yes, one or two times (21.64%)
- 3 – Yes, between three and ten times (14.93%)
- 4 – Yes, more than ten times (2.99%)
- 5 – Other (free text) (3.73%)

1.47% did not select an option [?]

Almost two-fifths of participants reported they had avoided reporting an incident or taking part in a discussion on AN/I in the last 12 months, because they were afraid of retributions of any kind.

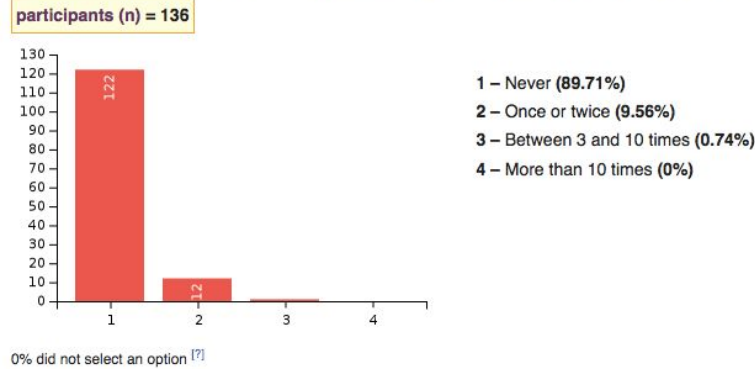
Almost two-fifths of participants reported they had avoided reporting an incident or taking part in a discussion on AN/I in the last 12 months, because they were afraid of retributions of any kind.

There was a significant split between responses of men and women was in answering the question “have you avoided reporting an incident or taking part in a discussion on AN/I in the last 12 months, because you were afraid of retributions of any kind?” 69% of women have hesitated to weigh in or report for fear of some form of retribution, with around 40% of men hesitating to weigh in or report for fear of some form of retribution.

Admonished and Non-admonished Editors:

In order to determine if there was a relationship between editors' willingness to get involved in discussions and their prior experiences in AN/I, participants were questioned about having been admonished in the last 12 months.

In the last 12 months, how often have you been admonished or sanctioned as a result of being involved in an incident that was reported to AN/I?



Almost 90% of survey participants have not been admonished or sanctioned as a result of being involved in an incident that was reported to AN/I in the last 12 months.

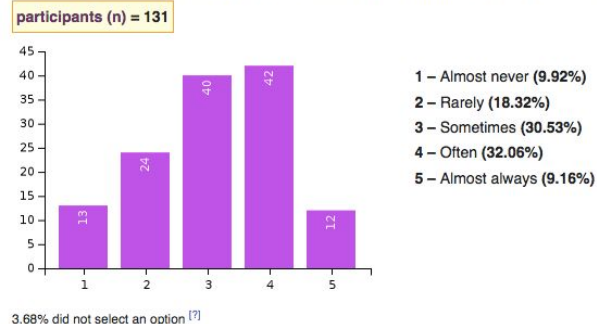
122 respondents have not been admonished, while only 14 have been admonished to varying degrees. The amount of experience editing seems to have no effect on whether an editor will be admonished or not. The more experienced contributors, those who have been editing for 5+ years, made up over 70% of participants who were admonished, as well as over 70% of participants who were not admonished.

The Effectiveness of AN/I Discussions and Reporting:

A large part of the quantitative questions in survey focused on gauging the effectiveness of the AN/I noticeboard in handling specific kinds of cases, which cases AN/I handled best and then how discussions on how AN/I cases were handled.

Threads typically stay on AN/I for the appropriate length of time.

How often, in general, would you say the following statements apply to discussion on AN/I? – Threads typically stay on AN/I for the appropriate length of time.



Four in ten of those who responded said that the threads "often" or "always" stay on AN/I for the appropriate length of time; three in ten instead said this only happened "sometimes".


Four in ten of those who responded said that the threads "often" or "always" stay on AN/I for the appropriate length of time; three in ten instead said this only happened "sometimes".

Quality of the discussions on AN/I

People's opinion about the quality of the discussion might indicate their willingness to participate in the discussion. It also speaks to the perceived lack of effectiveness of the noticeboard to resolve issues. Additionally, a part of the survey covered how users view the effectiveness of reporting on AN/I.

How often, in general, would you say the following statements apply to discussion on AN/I? – Discussion is focused and neutral.

participants (n) = 132



2.94% did not select an option [?]


Almost half of respondents said that discussions on AN/I are "almost never" or "rarely" focused and neutral.

49% of the survey respondents say that the discussions on AN/I are "almost never" or "rarely" focused and neutral.

Type of Cases

The graph below illustrates how users ranked AN/I in regard to its ability to handle certain kinds of cases. This analysis into 'certain kinds of cases' that AN/I does/does not handle well is explored in more detail in the qualitative section of this report.

On a scale of 1 to 5, how well are specific types of problems are dealt with at AN/I?



- *The majority of respondents (58.91%) said they believe sock puppetry cases are handled well on AN/I.*
- *Almost half of respondents believe that AN/I reports of personal attacks are handled poorly. Less than 10% of respondents felt they are handled very well.*


- *Three-fifths of respondents believed that problems involving bots or automated scripts were handled "reasonably well" or "very well" at AN/I.*
- *Two-thirds of respondents said they believe that impersonation accounts are dealt with well at AN/I.*
- *More than half of respondents said they believe that long-term user disputes are handled "reasonably poorly" or "very poorly" at AN/I.*
- *Fewer than 30% of respondents said they believe that topic-related problems were handled "reasonably well" or "very well" at AN/I.*
- *Almost three-quarters of respondents said they thought that copyright violations were handled "reasonably well" or "very well" on AN/I.*

Improvements to AN/I

While there is not clear support for a particular way to change AN/I, when grouping the answers in favor of change into one group, there is a trend towards some form of change. However, when the the answers are looked at separately, the respondents are divided on what changes could be implemented or needed for AN/I reporting.

In your opinion, what changes should be made to how reports are made on AN/I? (feel free to choose more than one answer)

127 responses



The above graph shows potential changes that could be made to the AN/I reporting process.

Only 26% of respondents thought that AN/I was fine the way it was. The most popular answer was “the use of structured reports” at 59%, the next most popular answer was “multiple options for reporting, e.g a mix of public and private reporting” at 35%, and the third most popular answer was “other, see next question” at 32%.


KEY FINDINGS (QUALITATIVE):

Questions from the survey were designed to gauge potential changes that could be made to AN/I. While there is a lack of consensus about any particular change, results show that change is suggested by a significant portion of respondents. There needs to be further follow up research to decide what kinds of changes are needed and support by the community.

Possible Changes To See the AN/I Reporting Process?

What are some other changes you would like to see in the AN/I reporting process?

participants (n) = 74
response buckets^[a] = 103



62 participants (45.59%) did not enter a response for this question

- | | |
|--|---|
| 1 – Clerking and moderators (14.86%) | 13 – Grouping related (4.05%) |
| 2 – Closer adherence to policy (6.76%) | 14 – Other technical solution (2.7%) |
| 3 – Boomerang (5.41%) | 15 – Misc filters / limits (4.05%) |
| 4 – General rules (2.7%) | 16 – Structured forms (4.05%) |
| 5 – Sanctions against abusers (5.41%) | 17 – Time limit / subpages (8.11%) |
| 6 – Fairness / Merit-based process (6.76%) | 18 – More structure in general (10.81%) |
| 7 – Punish bogus reports (4.05%) | 19 – Generate data / metrics (2.7%) |
| 8 – Referring people to different processes (2.7%) | 20 – General reporting (10.81%) |
| 9 – Training (4.05%) | 21 – Separate or ban uninvolved editors / non-admins (14.86%) |
| 10 – Misc. policy suggestions (4.05%) | 22 – Negative feedback (2.7%) |
| 11 – Environment (6.76%) | 23 – Unsorted (4.05%) |
| 12 – More civility in general (6.76%) | |

Written in feedback stresses a lack of consensus about a particular change but does seem to highlight that change is needed.

A user wrote in, “Some means of limiting non-involved parties needs to be found. At AN/I, there is a mix of users: (1) Those reporting. (2) Those attempting to helpfully respond and as needed to defuse potentially difficult situations. (3) Those who are trying in good faith to help, but who have an unfortunate habit of causing extraneous drama. (4) Those who have an axe to grind.

I have no idea how to restrict types three and four while encouraging types one and two - especially given that different people may have different interpretations of how to quantify efforts of others into those types.”

Another user said, “It’s not really “fine the way it is”, but other ways of doing it could potentially end up even worse. Private reporting sounds good at first, but it could turn into a so-called “star chamber”. Forms sound good at first, too, but not every situation fits into a form. If forms were optional, that might work.”

Another user wrote, “I would like to see more training/triage/ombudsman like work in filtering and supporting ANI reports; because of the rapid nature of the forum “I see a problem, folks let’s fix it!”


It tends to become a space for accusals, rather than support for feedback or soliciting constructive responses.”

What Do You Like About the AN/I Process:

One questions focused on what users liked about the AN/I process. In total, there were 66 write in responses with nearly a quarter of all responses being negative (“nothing”, “very little,” “not much. It’s a drama board mostly,” etc).

What do you like about the AN/I process?

participants (n) = 66
response buckets^[a] = 76



Common themes of what is liked about the AN/I process touched specifically on the open design of the board. Transparency and openness were the major takeaways on what is liked about the AN/I process: on how having an open board with an open and ad hoc process functionally allows for community input and transparency into how cases are handled and potentially solved. Additionally, the open design of the board allows for a quickness into solving, reading and weighing in on cases.

One user wrote, “I like its open nature: everyone with relevant information has the opportunity to provide it. Also, when it works, it works well. And sometimes it works perfectly: a succinct, informative report is made, an admin reads it, investigates, and takes action. That’s relatively rare, however...”

Another user commented on how AN/I “[is] easy, well documented, and consistent.” Allowing the community to weigh in publicly creates a system of community consensus, verifiability and transparency. But, even with its benefits, there’s an acknowledgment that the system is not perfect, but it’s what works at the moment.

AN/I serves a specific, necessary, and much needed purpose, but users acknowledge it’s not a perfectly designed system.

From a user, “...‘it may not be the best, but it’s the best we’ve got.’ When my reports did get answered, I usually was happy with the outcome, though it takes a lot of patience. It’s an ugly process to be involved in, though that’s by its definition. We need to get our hands dirty from time


to time to deal with such problems. Whatever process we have to replace this hypothetically, will be unpleasant...”

Which Type of Problems Does AN/I Work Well:

When asked “for which type of problems does AN/I work well,” there were 98 write in responses. Responses were focused on specific kinds of problems, which were sorted into three major groups: the complexity of the problems, technical issues or problems involving editor conduct.

For which types of problems does AN/I work well?

participants (n) = 98
response buckets^[a] = 166



38 participants (27.94%) did not enter a response for this question

AN/I seems to work best for more ‘obvious’ or less complex problems like sock puppetry, “short term” problems (problems that can quickly or easily be resolved), copyright violations, and ‘bot or automated script’ problems. However, complex versions of these problems, like a complex sock puppeting cases or a very complicated copyright violation may be handled better on another board. This seems to be reflected in the write-in answers, as well.

One user wrote, “AN/I works great for clear-cut rule violations, bot problems, etc. It does black-and-white very well.”

Another described the issues AN/I handles well as “housekeeping and basic admin stuff, like IP vandalism/sockpuppetry, blatant trolling, etc.”

Specifically, vandalism, bots, and sockpuppeting were mentioned frequently. AN/I seems to be able to handle cases that are best described as ‘simple or straightforward problems’, ‘technical problems’, ‘rules’ problems or ‘editor conduct problems.’ In their responses, some survey participants did not include specific problems that AN/I is able to handle well; rather, they spoke about more general concepts. Generally, the descriptions and suggestions were that AN/I handled cases best that were “objective and obvious cases” and “clear cut ones.” It seems that straightforward cases are ones that are easy to identify or easy to solve without a lot of administrator(s) input.

Examples of answers that were labeled as ‘straightforward cases’:

A user wrote in, “obvious instances of people breaking policies and guidelines.”

A user wrote in, “relatively simple issues that just need one admin or don't really involve disputes; disputes where one editor is clearly the disruptive one and can be sanctioned.”


A user wrote, “AN/I works great for clear-cut rule violations, bot problems, etc. It does black-and-white very well.”

What Does AN/I Not Do Well:

This question had 99 write-in responses. Some of the responses have been grouped into multiple buckets since the answers fit multiple kinds of labels.

For which types of problems does AN/I NOT work well?

participants (n) = 99
response buckets^[a] = 186



37 participants (27.21%) did not enter a response for this question

- | | |
|---|---|
| 1 – Conflicts with long history (21.21%) | 10 – Bad faith reports/"gaming the system" (5.05%) |
| 2 – Complex conflicts (24.24%) | 11 – Conflicts between factions (cabals) (13.13%) |
| 3 – Content conflicts (17.17%) | 12 – Conflicts with a certain type of editor (25.25%) |
| 4 – Certain topics (4.04%) | 13 – Interpersonal disputes (21.21%) |
| 5 – Unclear issues (11.11%) | 14 – Threats (1.01%) |
| 6 – Conflicts better suited elsewhere (7.07%) | 15 – Civility violations (18.18%) |
| 7 – Conflicts that attract attention (3.03%) | 16 – Reports from minorities, minority editors involved (3.03%) |
| 8 – Bot problems (2.02%) | 17 – (Almost) all problems (7.07%) |
| 9 – Admin conduct (4.04%) | |

The more complex, nuanced, complicated, or personal the issues, the more participants seemed to stress AN/I is not the right space to handle those problems. A complex problem may be something where the rules are ambiguous or contradictory, making solving a problem difficult, whereas a complicated problem could be a long running dispute between editors. Examples of these problems were described as cases that involved issues such as controversial topics, or issues involving civility problems and personal threats.

One user wrote in, “What AN/I doesn't do well is grey areas, long-term abuse, and taking action against "competence is required" editors whose problems are less clear cut.”

Another user highlights this 'grey area' problem by saying, "Ambiguous issues, issues where the topic is ideologically divisive, or where one of the users is new and the other one is a 'old hat' with plenty of friends (rarely is the discussion unbiased in these cases)."

The more ambiguous and/or more investigative work required, the harder it is to solve. Longer term problems that require more back and forth were viewed by admins as harder to solve, and something that AN/I didn't solve well. One user suggests that longer term problems would be better solved on another board or space such as "AN or ArbCom."

A user wrote in, "Anything where looking at large amounts of evidence is needed, where the outcome isn't obvious or where one or more editors are long-term so have supporters and opposers."

Another user wrote, "Complex interpersonal disputes, content disputes that have been ongoing for long enough or involve enough editors to have turned into conduct problems, complaints related to "incivility", any complaint whose specific circumstances touch on broader wiki-political issues and which therefore encourages a lot of soapboxing and off-topic personal commentary from observers."


Similar answers appeared like "anything personal," "personal disputes and personal attacks", "subjective cases (e.g. what is a personal attack)," etc. Lastly, feedback was also centralized on specific editors and admins who weigh in on AN/I. Users wrote about disputes that exist amongst or between certain editors or with "certain kinds of editors." This was a phrase mentioned repeatedly in the written in feedback to describe a certain kind of action that exists within AN/I. "Certain kinds of editors" describes the idea that some editors can "get away" with poor behavior which is looked over by administrators, and that is reflected within interactions on AN/I.

One user wrote, "Disputes with established editors who have built little fiefdoms for themselves based on intimidation, who throw their weight around, bullying and denigrating anyone who disagrees with them. They get away with it because no one is monitoring how often the "little people" have a problem with these established editors that think they can do whatever they please--which as it turns out is pretty much true. They do full reverts, don't explain, call other people's work "garbage" and the other person will be the one to get in trouble for "edit warring" or disruptive behavior or whatever. Multiple editors on an article where I was in conflict with one of these established bullies said they could see the justification for my complaints, but to a man, everyone of them told me I would not win anyway. He was established and knew how to manipulate Wiki's system and I was new and that would count against me. It didn't matter what was right or true, it only mattered that Wiki knew him and didn't know me."

Suggested Reasons that Incidents Would Not Be Handled Appropriately:

The survey asked a follow up, write in question to discover why participants think AN/I incidents would not be handled appropriately.

Have you avoided reporting one or more incidents to AN/I in the last 12 months, because you did not think it would be handled appropriately there?
 > If you did not select "no" above, why did you think those incidents would not be handled appropriately?



74 participants (54.41%) did not enter a response for this question

- | | |
|---------------------------------------|--|
| 1 – Avoiding drama (8.06%) | 9 – Boomerang effect (12.90%) |
| 2 – Complex issues (12.90%) | 10 – Easier to ignore the problem (8.06%) |
| 3 – Toxicity (12.90%) | 11 – Ineffective / inconsistent (8.06%) |
| 4 – Civility issues / threats (6.45%) | 12 – No chance of action (8.06%) |
| 5 – Biased participants (9.68%) | 13 – Bad past experience (6.45%) |
| 6 – Defensive cliques (19.35%) | 14 – Too difficult / too much scrutiny (3.23%) |
| 7 – Admin confidence (3.23%) | 15 – Better options (6.45%) |
| 8 – Certain users protected (4.84%) | 16 – Other / meta (3.23%) |

There were 62 write in responses. The write in answers were grouped into three major categories of responses: environment, reporting/process and participants. Generally, drama, bias, and toxicity or civility issues seemed to be some of the bigger deterrents that kept users from weighing in on AN/I. Drama, bias, toxicity, etc could be from the AN/I process itself, or the community that seems to allow specific participants to create these environmental deterrents. As well as the possibility of ‘boomerangs’ factor was mentioned.

One user wrote, “There is the boomerang factor. Who knows what I have in my history that people will comb through, so I put up with harassment or bullying or vandalism because the general rule is that boomerangs are allowed since there is no scope rule with AN/I reports”

Another user mentioned bias specifically, “It depends on who shows up; bias comes into play, e.g. people take sides based on whether they personally agree with the actions of people, not if people's actions follow policies and guidelines”

Groups of users and bad editor behavior is mentioned frequently, from edit count of those reporting being taken into account, to bias towards specific editors.

One user wrote, “Editors give far too much slack to people with whose ideas they agree, even when they admit the reported party has misbehaved, while at other times people are scapegoated without any evidence in the form of diffs being given--just simple assertions the "yes, X has been a problem for a long time, it's about time to ban X" with no demonstration of actual bad behavior.”

Another user said, “Certain users have a habit of bulldozing others.”

Specifically, how harsh, painful, and particularly ‘toxic’ the environment was of AN/I was cited as reasons why users avoided weighing in or reporting to AN/I.

A user wrote, "There are user cliques who will turn out to defend terrible, net-negative users. I am not a conspiracy theorist; I'm an admin with over a decade of service. This is simply true."

Another user said, "Because it is like a public hanging."


One user described this scenario, "There has been a significant decline in the standards of behavior accepted on Wiki and AN/I is where that decline is accelerated. For example, there was an edit war a while ago where someone was repeatedly putting "You smell like shit" on their user page. The AN/I discussion was dismissed with "Again with the nannying of the user pages." Experiences like that tell me anything goes."

These answers illuminate noteworthy hesitations for why the environment is seen as toxic and harsh and also provide insights into the perceived norms of AN/I. If editors and administrators are viewing AN/I as a place that cannot handle the nuances and labor of dealing with all kinds of conflict mitigation and harassment cases in a safe and conducive way, then users will not use AN/I as a place to report.

If You Could Change One Thing About AN/I, What Would it Be:

If you could change one thing about AN/I, what would it be?

participants (n) = 110
response buckets^[a] = 134



26 participants (19.12%) did not enter a response for this question

- | | |
|---|---|
| 1 – Clerking and moderators (21.82%) | 13 – Structured forms (5.45%) |
| 2 – Closing cases (4.55%) | 14 – Subpages and length (4.55%) |
| 3 – Boomerang (2.73%) | 15 – More structure in general (9.09%) |
| 4 – Misc rules (0.91%) | 16 – Private reporting (1.82%) |
| 5 – Fairness / Merit-based process (10%) | 17 – Generate data / metrics (0.91%) |
| 6 – Punish bogus reports (3.64%) | 18 – Misc reporting (2.73%) |
| 7 – Referring people to different processes (5.45%) | 19 – Separate or ban uninvolved editors / non-admins (10.91%) |
| 8 – Misc. policy suggestions (2.73%) | 20 – Punish derailers (5.45%) |
| 9 – Punish incivility (8.18%) | 21 – No change (2.73%) |
| 10 – Environment (5.45%) | 22 – Get rid of it (3.64%) |
| 11 – More civility in general (1.82%) | 23 – Not sure (2.73%) |
| 12 – Other technical solution (3.64%) | 24 – Unsorted (0.91%) |

This question had 110 write in responses that can be divided into two main categories: **technical structure** and **policy**. Within those two large groups of technical and policy structure, answers were then

coded into smaller groups like moderators/clerks, templates, general structure, new policy, filters/UI, forms or subpages.

Answers grouped into “technical structure” included feedback such as suggesting more structured data, which could mean creating specific forms or templates for to better structure and organize data. Users wrote in suggestions like “more structure,” “standardize reporting, to avoid walls of text,” “structured reporting (with opt-out allowed)” and “use subpages for issues to avoid edit conflicts.”

Answers grouped into “policy” were feedback that included more clearly defined rules, better rules for how cases are tried, clerking/moderation and noting a mixture of needing some private and some public reporting. Users wrote in suggestions like “Clerking to keep discussions focused” or “Enforce the civility guideline. We should not hesitate to sanction combative and disparaging behavior just because it happens to fall below the threshold of personal attacks.” Some feedback was helpful but very general such as, “instructions, forms, examples, admins who have time to investigate.”

Other specific write in suggestions were, “Allow users to make reports that are private or only visible to administrators.” Some feedback was more more general such as suggesting change, but not a specific kind, such “[change] the way it is structured. How AN/I is currently built is extremely conducive for agenda driven editors to watch the page all the time and continue drama and other "power maneuvers" against "other factions. I think having some fundamental changes in design could incentivise people to be civil and/or look for peaceful resolution, rather than always try to escalate.”

Responses mentioning ‘moderators/clerks’ was thematically one of the largest clusters. In that feedback, there were also detailed requests for administrators to take a more proactive role inside of AN/I to focus discussions, better mediate, and assist in archiving and ending cases. Currently, AN/I is an open board and adhoc space for editors to receive help, but there is no official requirement for administrators who give advice to remain involved until the case is resolved. This lack of clarity and structure is a part of the AN/I design as an open board, but can lead to confusion due to the freeform nature of discussion.

One user wrote: “more participation in long-running discussions by admins -- too often it appears that some admins avoid AN/I as a "dramah board", and the lack of admin oversight allows discussions to go on for too long without closure. Some method should be found to encourage admins -- perhaps by rewarding them in some way -- to deal with difficult AN/I complaints. Too often, discussions simply peter out as everyone loses interest, or lack anything new to say, but a discussion which has attracted many responses is obviously one in which the Wikipedia community has an interest in settling, and admins need to bite the bullet and wade in an make a decision. Overall, AN/I works decently well for all its faults, but I think the lack of admin participation is its biggest problem.”

The response above helps illustrate the community’s view on the challenges in how AN/I operates today. It also shows a desire to improve the process, recognizing AN/I’s (albeit imperfect) function.

CONCLUSION AND NEXT STEPS:


Overall, responses to this survey highlighted the general dissatisfaction with AN/I. However, the sentiment was often coupled with an acute awareness that AN/I can solve some problems or can be a necessary stage in the conflict resolution process. The qualitative and quantitative feedback shows AN/I handles

‘some cases’ well – such as problems that are easier to identify, or are straightforward or short-term – but complicated issues – like interpersonal disputes, controversial topics, or long term problems – are not handled well.

The qualitative responses collected best captures the nuanced feedback that users had for being dissatisfied with AN/I. Often the dissatisfaction was caveated with why AN/I does do something things well or how AN/I ‘may not be the best but it’s the best we’ve got’. In that sense, AN/I serves an important purpose, it can help mitigate cases of harassment early on, and helps keep cases out of ArbCom, which only handles very serious cases of harassment.

AN/I handles early forms of conflict mitigation well, such as straightforward cases, easier to identify or more simple problems, like shorter term problems, sockpuppetry, etc. This is reflected in the graph below.

On a scale of 1 to 5, how well are specific types of problems are dealt with at AN/I?



**from the quantitative part of the survey*

Suggestions for improvements included a desire for more participation and oversight from administrators, as well as new approaches to de-escalation and mediation. Respondents also highlighted issues with uninvolved participants on the noticeboard; some contributions from participants uninvolved with the cases were seen as unhelpful or distracting. Written in feedback from the survey suggested participation from administrators, creating moderators or assigning specific administrators to cases. This feedback could potentially help quell the issue of uninvolved participants weighing in on cases.

Another issue was the perception of ‘drama.’ A user wrote in specifically addressing how incidents are handled on AN/I, “because there is a lot of editors, admin and otherwise, who go to AN/I with an agenda, especially drama. Unbiased discussion on topic is hard to happen in any such scenarios.” Could assigning administrators to cases or designating a certain kind of moderator to alleviate this problem to help keep discussions focused and more effectively resolve cases? This is something that could be discussed by the community as a potential solution.

Additionally, while there was not concurrence about a specific type of structure, there was strong feedback requesting various forms of structure: structure in general, structural improvements like templates, structure related to policy (moderation or more involved admins), technical structure like use of subpages and structural improvements like filtering and UI.

Moving forward, Wikimedia Foundation will be following up and researching potential improvements based on this feedback. A future survey could research which types of cases are best to be solved publicly or to be solved privately. More follow up research is needed to further explore gender patterns in findings.

The Wikimedia Foundation realizes and respects that AN/I is community-led process; any potential improvements will need to be subject to community consensus and implementation.

Next steps for the Anti-Harassment Team and the Support and Safety team involve building out product takeaways and hypotheses from the survey, continuing administrator and editor interviews specifically on conflict mitigation and reporting harassment, a product audit of the different spaces, places and tools for reporting on Wikimedia projects, and more general fact finding on reporting systems. All interested people are invited to provide feedback. Please post questions and comments on the talk page.

APPENDIX

Detailed analysis with tables and charts:

Admonished and Non-admonished Editors:

In order to determine if there was a relationship between editors' willingness to get involved in discussions and their prior experiences in AN/I, participants were asked how many of them had been admonished at some point by administrators in AN/I cases. There were a low number of admonished participants compared to the number of non admonished participants. Even with that low number, the group was worth investigating to compare user patterns and trends to the non admonished group. 122 respondents have not been admonished, while only 14 have been admonished to varying degrees. The amount of experience editing seems to have no effect on whether an editor will be admonished or not. The more experienced contributors, those who have been editing for 5+ years, made up over 70% of participants who were admonished, as well as over 70% of participants who were not admonished.

Both admonished and non-admonished groups are reporting incidents to AN/I. Out of the non admonished group, only 34% have never reported, with 48% filing incident reports once or twice a year. Out of the admonished group, only 21% have never reported, with 29% filing a report once or twice a year and 43% filing between 3-10 times a year.

How Often Have you Reported Incidents to AN/I in the last 12 months?

Amount	Not Admonished	Admonished
Never	33.6%	21.4%
1-2	48.4%	28.6%
Between 3-10	15.6%	42.9%
10x	2.5%	7.1%

Both admonished and non admonished users are weighing in uninvolved cases; 76% of not admonished users weigh in on cases (frequency varying) with 86% of admonished users weighing in cases they are not involved in.

How Often Have You Been an Uninvolved Participant in discussions on AN/I in the last 12 months?

Amount	Not Admonished	Admonished
Never	24%	14.3%
1-2	27.3%	14.3%
Between 3-10	28.1%	35.7%
10x	20.7%	35.7%

How Often Have You Been Involved in an Incident Reported on AN/I in the last 12 months?

Amount	Not Admonished	Admonished
Never	31.1%	0
1-2	45.9%	50% (7ppl)
Between 3-10	17.2%	50%
10x	5.7%	0

Both admonished and non-admonished visit AN/I to read and follow reports that do not involve them, with 94.4% of not admonished visiting AN/I to read or follow reports to some degree, and 86% of admonished visiting AN/I to read or follow reports to some degree.

How often do you visit AN/I to follow or read reports that do not involve you?

Amount	Not Admonished	Admonished
Never	6.6%	14.3%
1-2	10.7%	0
Between 3-10	20.7%	42.9%
10x	62%	42.9%

Editors who have been admonished were only slightly more active in AN/I in comparison to those who have not. Those that have been admonished do continue to weigh in on AN/I reports, and file reports themselves. The amount a user uses AN/I and a user's knowledge or expertise with AN/I does not seem to affect if a user will be admonished or not.

Gender Patterns:

Women respondents in the survey reported incidents slightly less than men with 44% of women saying they never report, 50% reporting once or twice, and 6.3% reporting between 3-10 times. Men report somewhat more, with 34% remarking they never report, 43% reporting once or twice, 20% reporting 3-10 times, and 3% reporting more than 10 times.

Generally, women are visiting AN/I to read or follow incidents. Around 8% of male participants said they never follow reports that don't involve them; all of the 16 female participants polled read or follow reports that do not involve them to some degree. 75% of the female participants are reading or following more than ten reports that do not involve them compared to 61% of male participants. However, just because women are reading reports, that does not mean they are weighing in. 81% of women participants are uninvolved participants in discussions on AN/I, with 75% of men being uninvolved participants in discussions.

Additionally, another finding that did emerge that needs further studying, research, and analysis: namely, that none of the female participants in the survey had been admonished, while almost 10% of the male participants had. Could this suggest that men who get admonished stay on the AN/I, while women who get admonished leave? More follow up and research is needed to further explore these findings.

Detailed Methodology

Qualitative section

The qualitative part of the survey had six questions with write in answers. The questions asked participants for thoughts and feedback on statements like "if you could change what thing about AN/I, what would you change" or "what does AN/I do well/not do well?" The write in answers were then sorted into different labeled groups of feedback, such as clerking/moderation, enforcing copyright violations or legal threats, filtering, structure, UI suggestions, templates, editor conduct, and general requests for more structure. Some write in answers were grouped into multiple labels since the feedback fell into multiple groups. At a high level, the answers generally stemmed into two kinds of feedback: policy and product feedback, which could help guide new kinds of policy or product to potentially be implemented within AN/I. Even when answers were gauged more towards users' feelings towards AN/I such as if AN/I was liked or disliked, the feedback could still often be sorted into groups that focused on policy or product implementations.

Each question had its answers grouped separately, though some labels were used for grouping answers for different questions. For example, some form of 'functional' or 'technical' feedback were used for abeling across all questions whereas only one question had a group of responses that were labeled as 'policy.'