IP Masking Impact Report

Claudia Lo | 2019-07-22

Contents

Introduction

Currently, IP addresses are captured on account creation and on any edit action (including participation in structured discussions e.g. Flow). IP addresses associated with registered users are available only to a select group of users. However, the IP addresses for unregistered users are displayed and archived, publicly, on an article's edit history page.

This document will refer to users who have created a Wikipedia account, and are using it to edit Wikipedia, as *registered users*. Users who are not logged in when they interact with Wikipedia, and therefore have their IPs captured and publicly logged on edits, are *unregistered*

users. While other terminology exists, including "anonymous users" and "IP users", I have chosen this terminology because it suggests less of a mutually exclusive split between users editing with a Wikipedia account, and a user editing while not logged into an account, while being more descriptive than either.

Unregistered users' IP addresses have been captured and published on MediaWiki since its earliest incarnations. This ubiquity means that public IP addresses are used in many processes, broadly separable into governance and social processes. Governance processes are those that are involved in the day-to-day running and maintenance of our projects, while social processes are those revolving around community norms and values, and ongoing debates about the place and treatment of unregistered users.

Glossary

Internet Protocol address (IP, IP address)

A numerical address assigned to each device that uses the Internet Protocol for communication; for the purposes of this paper, this is any device that connects to the Internet. IP addresses are unique sets of numbers, and generally assigned by the internet service provider (for residential use) or organization to which the device belongs (for commercial or institutional use). There are two main types of IP addresses in common use, IPv4 and IPv6. However, for the purposes of this document, the technical details of handling one over the other are not necessary. The devices are often routers rather than computers themselves, or mobile devices if they are on a data network.

Because of the way IP addresses are assigned, they can be used to locate devices geographically. Additional information associated with particular IP addresses can also be viewed via a WHOIS lookup. Therefore, IP addresses may be considered identifying information, depending on the practices of the ISP assigning the address, and any extra information associated with it. Particularly sensitive IP addresses or ranges, such as those

belonging to the United States Senate or the Parliament of the United Kingdom, are listed on administrators' pages to remind them not to block these addresses.

Static IP

A static IP address is one that does not change after assignation. This may apply to residential or commercial IPs. A single static IP address may also be used for a fairly large institution, depending on how that network was set up.

Dynamic IP

Dynamic IP addresses are IP addresses that change. This means that a single device may be represented by a range of IPs, or that a single IP may come to be associated with multiple devices over time. Dynamic IPs tend to move between a specific range, with its size depending on the policies of the service provider.

IPs of mobile devices tend to be dynamic, again over a shared range used by many other mobile devices on the same service provider network. Even a seemingly-static IP address used by a mobile device is likely to be shared on many devices due to carrier practices. If these devices connect to WiFi networks, each new connection will use the IP address of the router instead. Suffice to say, the IP address for mobile devices can vary wildly, and one device can have more than one public IP address.

WHOIS lookup

A WHOIS check, or WHOIS lookup, is a query-and-response protocol that can be used to look up the registered information associated with a given IP. In practice, this means it functions as a tool that can provide the name of the organization that has registered the IP, where it is located geographically, and when the address was registered. Many WHOIS services exist online.

"Collateral"

The term “blocking collateral” refers to the cost of accidentally blocking uninvolved users as a result of an improperly-set block. This cited as a guiding principle behind limited use of range blocks in particular. It is very important when considering blocks on mobile devices, specifically because mobile devices tend to use many, shared, IPs. The term is often shortened to “collateral”. While other projects may have the same concept, this term is specifically used on English Wikipedia.

Sockpuppet

Sockpuppets refer to multiple accounts, operated by the same person, that are not marked as such. English Wikipedia policy allows for the use of alternate accounts in very specific situations, and all alternate accounts must be clearly marked as such and associated with a main account. Sockpuppets are therefore treated with suspicion due to their association with vandalism, astroturfing (using multiple accounts to give the false appearance of grassroots support) and other bad-faith behaviours, where the use of multiple accounts is an attempt to frustrate administrator action and obfuscate the identity of the main actor.

Technical evidence

This term refers to any information gained on the basis of technical information on a user: their IP address, user agent information, or other similar information. The fullest spectrum of technical evidence is currently only available to those with the CheckUser permission, since they can view both IP addresses and the associated user agents, or vice versa. Technical evidence is generally used in concert with behavioural evidence to assess whether or not two given accounts or unregistered IPs are being operated by the same person, or a closely-organized group of people.

Behavioural evidence

In contrast to technical evidence, behavioural evidence is information gathered on the basis of a user’s activity patterns. This includes patterns in when the user is active, writing styles, editing quirks, posting habits, username patterns, edit summary usage, and other such

behaviours. Behavioural evidence is open to any user who spends the time to collect this information, and is heavily weighed when trying to assess if there is a single person behind multiple accounts. The ability to quickly and accurately collect, analyze, and use behavioural evidence is behind the "duck test" principle, elaborated upon in WP:DUCK.

Key Takeaways

IP usage

Governance

- IP addresses are valuable as a semi-reliable partial identifier, which is not easily manipulated by their associated user.
- Depending on provider and device configuration, IP address information is not always accurate or precise;
    - IP information or technical knowledge is used to support additional information ("behavioural knowledge") where possible.
- Deep technical knowledge and fluency is needed to make best use of IP address information, though administrators are not currently required to demonstrate such fluency to have access.
- The information taken from IP addresses, such as address range, and institution type, as well as information derived from user agent such as device type, significantly impact the course of administrative action taken.

Social

- The issue of whether to allow unregistered users to edit has been a subject of extensive debate; so far, it has erred on the side of allowing unregistered users to edit.
- The debate is generally framed around a desire to halt vandalism, versus preserving the ability for pseudo-anonymous editing and lowering the barrier to edit.

- There are major communications issues when trying to talk to unregistered users, largely centered on lack of notifications, and there being no guarantee that the same person will read sequential messages sent to that account.
- There is a perception of bias against unregistered users as a consequence of their association with vandalism, which has persisted to algorithmic bias.

Impacts

- IP masking will significantly impact administrator workflows and may increase the burden on CheckUsers in the short term.
- In the short term, our administrators' ability to manage vandalism will be greatly hindered. This can be mitigated by providing tools with equivalent or greater functionality, but we should expect a transitional period marked by reduced administrator efficacy.
- We must be careful to preserve or provide alternatives to the following functions currently fulfilled by IP information:
    - Block efficacy and collateral estimation
    - Some way of surfacing similarities or patterns among unregistered users, such as geographic similarity, certain institutions (e.g. if edits are coming from a high school or university)
    - The ability to target specific groups of unregistered users
    - Location- or institution-specific actions (not necessarily blocks); for example, the ability to determine if edits are made from an open proxy, or public location like a school or public library.
- Depending on how we handle temporary accounts or identifiers for unregistered users, we may be able to improve communication to unregistered users.
- Underlying discussions and concerns around anonymous editing, anonymous vandalism, and bias against unregistered users are unlikely to significantly change if we mask IPs, provided we maintain the ability to edit projects while logged out.

Governance

Currently, the majority of MediaWiki's blocks work via IP, even if these IPs are not always surfaced to the admin performing the block. While IP addresses are taken as a reliable indicator of unique identity, the fundamental issue is that knowing the IP address associated with a pattern of editing behavior does not guarantee that there is a single individual performing those edits, nor that it is the *same* individual exhibiting this behavior. Therefore, IP addresses are used in a few key ways for governance on our projects:

- As a trackable identifier for edits made by unregistered users, against which administrative action can be taken,
- As one piece of evidence, amongst others, that can be used to establish whether or not the same individual is operating multiple accounts.

Blocks

Single IP block

Single-address blocks bar a single IP address from editing the site, or specific pages in the case of partial blocks, for a specified duration. These blocks are most effective on static IPs, generally a static residential or public address. They are less effective for dynamic IPs, or if the person being IP banned understands how to change IP address.

IP range block

An IP range block stops all IPs within a certain range from editing for the duration of the block. IP ranges for IPv4 and IPv6 ranges look different and have different implications, but the principle is the same. If a cluster of IPs are being used for misconduct, administrators can block a range of IP addresses that encompass this cluster, preventing whole groups of bad-faith actors or a single determined bad-faith actor from causing further harm to the project. This is handled via CIDR syntax, though this requires a high level of technical knowledge.

Currently, MediaWiki allows blocks no larger than /16 for IPv4, and /19 for IPv6; in essence, this is a range-block wide enough to cover every IP assigned by a single internet service provider (ISP), which some cases, might be enough to block an entire country. Administrators are expected to check the coverage of ranges they intend to block in order to assess collateral damage, and a gadget exists on English Wikipedia to warn administrators if they attempt to block a sensitive or very wide IP address or range.

Specialized IP range blocks

All IPs can be queried using a WHOIS lookup, which generally will return the name of the organization associated with the IP. In the case of business or institutional IPs, this is generally the name of the organization that uses that IP; for residential cases, it is often the name of the ISP that has assigned it.

Therefore, it is possible to determine when a cluster of IPs used for bad-faith activity comes from a single institution. This has led to the creation of certain specialized IP range blocks, expressed using templates.

School blocks

These blocks use Template:School block and Template:School block hard, though technically they are the same as IP range blocks. These range blocks are used when administrators believe that a set of IPs, originating from a school, are routinely engaged in vandalism. Administrators can determine this from the WHOIS lookup.

Wikimedia Foundation office block

There is a block placed on all Wikimedia Foundation office IPs, in order to prevent staff from editing without first logging in.

Dynamic IP concerns

IP range blocks have been occasionally used to counter vandalism or other bad-faith actions originating from dynamic IPs. Depending on the ISP's practice, residential IPs can be either static or dynamic over a small range. Mobile IPs are also often dynamic, depending on carrier practices. However, as mobile devices hop between networks over the course of the day, the IP address associated with the device will almost certainly change, making it a de-facto dynamic address.

If an administrator has knowledge of ISP practices, a range block can be used to stop bad-faith action coming from a dynamic residential IP. Range blocks are very rarely used on mobile ranges, since the risk of collateral damage is too high.

Range blocks for dynamic IPs are rarely used. Instead, autoblocks, which automatically apply a cookie-based block, are used to address the issue, since they are very effective at obstructing low-effort bad-faith behaviour without requiring the level of IP knowledge that a range block would need. As mentioned in the glossary, mobile IP addresses in particular are ill suited for IP blocks because of the almost-inevitable collateral it would cause; since most mobile IPs are shared, and they are dynamic, any range block wide enough to cover the shifting IPs of a bad-faith mobile editor would be guaranteed to catch many other good-faith or uninvolved users. This method also does not require the depth of technical knowledge required to set a range block with the proper CIDR syntax.

Autoblock

Autoblocks are available on any project. It automatically blocks any IP used by a blocked and registered user. This is intended to combat IP hopping, where a blocked and registered user purposefully changes their IP address in an attempt to circumvent existing blocks. However, as is noted on the autoblock policy page, the potential for collateral is high if the initially-blocked account uses a dynamic IP address as a result of ISP policy, since any other user who is then assigned the same IP address down the line will be blocked from editing through no fault of their own.

Setting an autoblock is a configurable option for administrators to choose when setting blocks, and administrators cannot access the actual IP of the autoblocked user in this process. Autoblocks are handled entirely via MediaWiki. English Wikipedia has its own set of best practices for administrators to follow when using autoblocks, such as suggesting them in cases where it is likely that the blocked user will attempt to evade the block.

IP Block Workflow

A typical IP block incorporating WHOIS might go as follows.

- Receive report of vandalism from unregistered user;
- Run WHOIS on IP address to determine type (static, residential, commercial, mobile):
    - If static residential or commercial, single IP block;
    - If dynamic residential, rangeblock or block with autoblock;
    - If mobile, autoblock single IP, after checking for collateral;
    - If dynamic, but worried about collateral, consider autoblock (which applies a cookie to account blocks) or AbuseFilter;
    - If unsure, and vandalism is severe enough, hardblock (no edits at all from the range or IP, whether registered or not).
- Wait to see if block is effective and that type of vandalism stops occurring;
    - If it reoccurs, consider broadening the parameters of the block.

The duration of the block, and the width of the range, is determined by the severity of the vandalism, and the patterns exhibited by the bad faith user: if they change IPs, a rangeblock is more suited than a single IP block, but if the IPs they use are either shared or occupy a very wide range, the potential for collateral is high enough that another method is warranted.

Page protection

In order to limit edits or changes to certain pages, administrators can set different levels of page protection on an article. This restricts editing privileges for that page to limited groups of users, depending on the level of page protection.

For this project, the most important level of page protection is **semi-protect**. A semi-protected page cannot be edited by unregistered users, or non-autoconfirmed users. The exact definition for an autoconfirmed user varies by project, but generally covers very new accounts as defined by both low age and edit count.

Pending changes

Occasionally pitched as an alternative to page protection, pending changes places all proposed edits, from registered or unregistered users, in a queue for review by an editor with the appropriate permissions. On some projects, such as German Wikipedia, pending changes is turned on for all articles.

AbuseFilter

AbuseFilter is a MediaWiki extension that allows users with the appropriate permissions, usually administrators, to set specific actions to be taken on certain triggering conditions. AbuseFilter works with IPs and IP ranges. For example, to combat a vandal whose primary mode of vandalism is uploading and redirecting existing article images to inappropriate pictures, an administrator might set AbuseFilter to block all image uploads coming from a given IP or IP range.

CheckUser

CheckUser can refer either to a user with the CheckUser permission, or the CheckUser extension itself.

CheckUser tool

The CheckUser permission grants access to the extension of the same name. It is a tool that allows users to see:

- All IP addresses associated with a given user,
- All user agents associated with a given IP address or range,

- All edits associated with a given IP address or range, regardless of whether they were made by a registered or unregistered user.

The tool was initially released in February 2007, and has received few significant updates since. It was preceded by a tool called Espionage (sometimes "Userip"), created in 2005 by a community member. It presupposes a high level of knowledge and working fluency with IP addresses. A more detailed walkthrough of its features, including the information it returns, can be found in its MediaWiki page.

By custom, CheckUsers may not use this tool to "fish"; that is, they must have a clearly noted reason in order to use the tool. All uses of CheckUser are logged, and CheckUsers generally must provide a reason (often referring back to a specific case or incident) for their use of the tool. However, it is difficult to verify whether or not this level of logging and tagging takes place for every use of the tool.

CheckUser role

Approved CheckUser candidates must currently sign a confidentiality agreement with the Wikimedia Foundation before they are granted access. Very few users across the entire movement have access. Those that do have access are generally limited to CheckUser for a single project. Only ombudsmen, who oversee complaints about CheckUser abuse among other concerns, and Foundation staff members have global CheckUser permission. Notably, CheckUsers are not trained in the use of the tool specifically, nor do they have to demonstrate prior knowledge or facility with IP addresses. The implication is that, through regular administrator work and beyond, candidates are expected to have received a sort of vocational training with regards to IPs and what information they provide. CheckUsers generally also have the checkuser-log permission, allowing them to view the logs for uses of the tool. They also have access to CUwiki, a private wiki that is used as a record-keeping space as well as a place to communicate to other CheckUsers.

Uses of CheckUser

Currently, CheckUser is often used in cases where administrators suspect sockpuppets or alternate accounts are being used. CheckUsers may be called upon to verify if two accounts share an IP. Such a confirmation, in addition with other information such as similar posting patterns, writing styles, activity patterns, usernames or any other contextual information, is used to determine whether or not two accounts belong to the same person. The prohibition against using CheckUser to trawl, or in other words using it to start an investigation into an account's behavior, fundamentally means that IP addresses are rarely if ever used as the sole evidence proving that a single person is behind multiple suspected sockpuppets. CheckUser comes into play only after there is already suspicion directed at an account, not before.

Social

One important note to make is that, aside from the issues of communicating with an anonymous account, many of the issues highlighted here are not unique to unregistered users and apply to new users as well, to some degree. However, the interaction of these issues with the fact of an unstable social identity means that the overall experience and outcome is unique to unregistered users.

A non-exhaustive list of existing essays on the subject are listed below.

- WP:IPs are human too
- WP:IP addresses are not people
- WP:Why create an account?
- WP:Perennial proposals/Prohibit anonymous users from editing
- WP:Counter-vandalism unit/Vandalism studies
- WP:Not every IP is a vandal
- WP:Seigenthaler biography incident

Anonymity and anonymous editing

One interpretation of Wikipedia as the encyclopedia that "anyone can edit" is that anonymous users should also be allowed to edit. However, there are two key meanings of anonymity that can sometimes be conflated, confusing concerns of privacy with concerns about social identifiability.

In the context of online privacy or security, anonymity refers to a state in which an individual user using a given service reveals no public identifying information about themselves, and does not have any identifying information harvested and exposed publicly.

In the context of an online social community, anonymity refers to a state in which an individual user does not have a stable public-facing identity or identifier. Whether this identifier correlates to their offline identity or is used only in a single community is irrelevant.

A user may be anonymous in a privacy sense, yet still be identifiable in a social sense; that is to say, their online identity may not be connected with any other aspects of their identity outside of this community, yet still be recognizable as a single "voice". For example, a registered Wikipedia user may have set up an account solely for editing purposes and be cautious enough to never reveal identifying information about themselves, yet their long association with a single account name provides a stable social identity.

Conversely, a user may be anonymous in the social sense, yet be identifiable from a privacy perspective. For example, a user who edits Wikipedia using a computer at a public library exposes their geographic location since their IP address will be logged, yet there is no guarantee that future edits from that IP will come from that same person. Indeed, in such an example, there is no guarantee that any two edits from that IP address are being made by the same person. While there is a loose group of dedicated editors who refuse to register accounts, it is safe to say that they are in the minority of unregistered users.

Returning to the issue of anonymous editing, the ability to edit without requiring an account is one that has been contested over the history of our projects. Arguments for it include minimizing barriers to participation in our projects, and interpreting "everyone can edit" as allowing unattributed contribution. Arguments against it generally stem from a desire to restrict the ability of bad-faith contributors from accessing the projects, since a low barrier to participation also means it is easier for harmful, destructive, or simply ill-fitting contributions to be made.

Practically speaking, while no project has disallowed all unregistered user edits as a matter of course, unregistered users are generally restricted in what types of contributions they can make as compared to registered users. For example, unregistered users cannot start new articles or upload files. Furthermore, unregistered users' lack of a stable social identity makes it difficult for them to communicate and fully participate in their project's community in several ways.

Communication

Unregistered users have talk pages associated with their IP address. The presence of new messages on this talk page are marked with a bright yellow banner at the top, but they do not receive a notification through MediaWiki's notification system. Whether or not the unregistered user checks their talk page, therefore, relies in part on their existing familiarity with user talk pages, and their susceptibility to banner blindness.

Unregistered users cannot be pinged. Additionally, unregistered users cannot email other users. These aspects pose a significant barrier for communication.

While not being able to email other users is less critical, generally speaking, there is one important exception to this. Requests to delete sensitive information from a project, also called a revdel (revision delete) request, are meant to be done via email so that the request is not made on a public administrator's noticeboard. However, since unregistered users cannot send emails in MediaWiki, an unregistered user making a revdel request has little choice but to ask publicly.

This is a major issue for the security and privacy of the unregistered user, especially since revdel requests are typically made for very sensitive information or information about a minor. This lack of email functionality also limits an unregistered user's ability to contact Trust and Safety should that be required, although to a lesser extent, as they still have the option to send an email linking to the relevant pages using their personal off-wiki accounts.

User talk pages form the basis of most internal wiki communication, especially from experienced users trying to contact new users. The lack of a notification makes these important messages far more difficult to see. There is also little or no way to be sure that the person reading that IP address talk page is the intended addressee, especially if the IP address is associated with something like a public library computer, where multiple people may be editing from the same IP address with no contact between each other.

Because of the assumption that many of these unregistered editors are also new editors, experienced editors have created a slew of templates aimed at automating or streamlining repetitive messages. However, the clearly impersonal nature of these messages, combined with the existing issue of banner blindness, makes them easy to ignore even if the intended recipient happens to be reading them (which, again, is not guaranteed). Even though some of these templates, such as Template:Notavandal allude to the multiple users of a shared IP, this means very little to an editor who does not already understand or know what an IP address is, or why they might be sharing one with someone else.

Lastly, the obviously not purpose-chosen nature of IP addresses as identifiers has a dehumanizing effect; this can be inferred from the existence of such policy pages as WP:IPs are human too, which encourage registered users to treat unregistered users in good faith. And, while unregistered users may contribute in community discussions, they are generally barred from voting in administrator elections or other issues of project governance on English Wikipedia.

Vandalism

A 2007 study, conducted by a Wikipedia user, indicated that about 80% of vandalism comes from unregistered users, but that vandalism represented only about 20% of total edits made by unregistered users. While there have been few studies on the topic since, the perception that unregistered users are responsible for the majority of vandalism on Wikipedia fuels a perennial argument about whether or not Wikipedia should allow edits from unregistered users at all.

Accusations of bias

On the flip side of the vandalism argument, there is the counter-argument that unregistered users are subject to undue scrutiny as a result of the perception that unregistered users are more prone to vandalism, disruptive behavior, or have something to hide. This argument is

based both on the slow curtailing of permissions held by unregistered users, and the perception that administrator action against unregistered users are overly harsh.

This worry about bias has also extended to bias in algorithmic tools used by Wikipedians, such as in ORES. A recent Research Showcase discusses the issue of ORES disproportionately flagging edits made by unregistered users as low-quality, when compared to the human-labelled training data.

Impact on governance

Given the central position of IPs in blocking, we should expect that masking will significantly disrupt workflows and general administrator efficacy. Even if we provide tools with similar or greater functionality, there will be a transitional period where administrators have to learn how to best use them. Therefore, even in this best-case scenario, we must expect some short-term loss in our volunteers' ability to manage vandalism and other damaging acts on our projects.

There are two potential projected scenarios that will come about as a consequence of the IP masking project.

Scenario A: IP addresses are obscured for everyone except those with CheckUser and related permissions.

Scenario B: IP addresses are obscured for everyone, including those with CheckUser and related permissions.

In order to properly assess the potential impact of the project, I will mark scenario-specific obstacles; otherwise, assume that the potential issue is common to both scenarios.

Potential workflow disruption

Currently, WHOIS information is used as the basis of important decisions with regards to the appropriate administrative actions. Therefore, checking WHOIS forms an early step in administrator workflow. More precisely, the ability to spot patterns and similarities between actionable behaviour taken by unregistered users, part of which is made possible by viewing their IPs, is what is critical here.

For scenario A, I would imagine that a likely short-term outcome would be **increased CheckUser workloads**, as administrators adapt to the lack of public IP information by attempting to access it via the proxy of a CheckUser. However, as my investigations into English Wikipedia sockpuppet investigations (a common working site for CheckUsers)

indicates, there are relatively few CheckUsers and they are already under strain as is. This is evidenced by the backlog of sockpuppet cases, and the relatively strict rules around when a CheckUser will actually run a check. In the long term, we might optimistically hope that administrators will adapt to the loss of IP information by adapting their processes to compensate. This compensation could come in the form of greater reliance on administrator intuition with regards to behavioural evidence, or the development of alternative technical identifiers.

In scenario B, the disruption would likely be much greater, especially since the distinguishing ability of CheckUsers would be essentially removed. Depending on the form that the IP masking will take, there may be information accessible only to CheckUsers that regular administrators do not have, due to the CheckUsers' ability to cross-reference masked IPs against user agents. Without knowing the specific technical solution we will implement to mask IPs, it is hard to say more.

Key functions to preserve

Impact assessment

WHOIS information allows administrators to know what type of IP address an unregistered user employs. In turn, this information allows them to assess potential collateral. A rangeblock placed on a limited residential dynamic IP will impact far fewer people than a rangeblock placed on a mobile IP range. Similarly, it also allows them to assess the efficacy of a block. A single-address block on a static IP will be harder to evade than a single-address block on a dynamic range.

Therefore, we need to provide some method for administrators to assess potential collateral, or other possible undesirable outcomes, when placing a block or taking other administrative action. Relatedly, administrators need some method to estimate the efficacy of a given administrative action.

Precision of action

For this, I am assuming that one outcome of IP masking will be the need to adopt a new identifier for unregistered users, possibly unrelated to their IP. The primary method of limiting or precisely targeting a certain section of unregistered users, by using IP-specific targeting, will need to be replaced by something else.

For example, AbuseFilter can currently be used to block specific actions carried out from within a very specific IP or IP range. In the absence of having easily-accessible IP addresses, we will need to provide some way to target a specific unregistered user, or groups of unregistered users.

Categorization of users

Related to the above issue of precision, we need some way of being able to surface patterns and similarities in the actions of both registered and unregistered users. Currently, IPs serve this purpose by providing limited geographic and institutional information that requires a degree of technical know-how to spoof. This information, whether accessible publicly or via CheckUser, allows administrators to note down patterns that allow them to identify long-term vandals, and help confirm or deny the presence of sockpuppets, to name a few functions. For example, it is highly unlikely a single person will rapidly switch countries or continents, and so this can be a hallmark of either proxy use, or an account being hacked and used by someone other than its owner.

This also allows administrators to avoid placing blocks on users who perform edits that look similar to multiple account abuse, such as student work, or edit-a-thons, based on the location

Therefore, we ideally want some way of surfacing connections or similar patterns between groups of users in a way that does not involve revealing IPs, or the private information that masking IPs is meant to protect. This should also allow administrative action to be taken against these surfaced connections or groups of unregistered users.

Location- or IP-based action

Some administrative actions placed on IPs are done largely on the basis of what those IPs represent, not what editors using that IP have done. For example, school blocks are a specific type of block meant to curb vandalism of a narrow set of pages, while keeping open the option of allowing younger students to contribute to Wikipedia. Knowing that an IP address belongs to a highly-trafficked, publicly accessible device (such as a library computer) can be useful for administrators managing its use or trying to communicate with those unregistered users. Since proxy use is not permitted on Wikipedia, all IP addresses used as open proxies are blocked on that basis. Lastly, there is a block placed on all WMF office IPs, to ensure that staff log in before editing. Additionally, sensitive IPs, such as the ranges used by the US Congress, are known to administrators, and they use this information to treat unregistered or logged-out edits from those ranges differently.

Therefore, in these limited cases, we may want to surface limited geographic information to allow our administrators to act based on the place from which those edits originate.

Unified global identifier for unregistered users

One way in which attributing edits to IP addresses is useful is that it is a stable globally-consistent identifier, allowing functionaries to check for cross-wiki vandalism. In the absence of IPs, any temporary or anonymous username we create should keep this function, to facilitate cross-wiki abuse mitigation.

Impact on social processes

By comparison, masking IP addresses may have a less disruptive impact on social processes than governance processes. Our major consideration is in how unregistered or logged-out usernames will be displayed instead of showing their IP address.

Anonymous editing

The IP masking project is very unlikely to change the conversation around anonymous editing, anonymous vandals, or right to privacy. So long as the ability to edit while not logged in remains, the underlying concerns behind these issues will also remain.

Communications

Currently, communicating to an unregistered user is a fairly fraught endeavour, especially since rightfully speaking, communications are grouped by IP rather than the user *per se*. Depending on the technical solution we choose to come up with usernames for unregistered or logged out users that does not surface their IP, this is an opportunity to improve communications with logged out users. Again, depending on the implementation, we may have an opportunity to create a temporary identity for unregistered or logged-out users that stands a better chance of delivering that message to the person behind the screen. However, details will remain out of reach until work begins on technical implementation.