Keeping Watch Over the Fediverse: Mass Surveillance in Non-Centralized Social Media

Eric Fassbender

February 17, 2025

Abstract

Non-centralized social media appears to be undergoing a “Killer Hype Cycle”, where many users dissatisfied with centralized corporate platforms are identifying Mastodon as an alternative. With the influx of users there has been an increase in available data for researchers of social media and communication. Much of this work has focused on the everyday end user: someone who is using the platform to share personal information or consume media. However, these platforms have other uses. Corporations seek to mine federated media for their own endeavors, and state agents catalog the information for various uses. In this paper, I analyze data from the BlueLeaks dataset, 270 gigabytes of leaked messages, training, inventory lists, and membership logs from U.S. Fusion Centers, to determine state surveillance intrusion into non-centralized platforms. By viewing this data through the model of Abstraction, Topology, and Scale proposed by Zulli, Liu, and Gehl, I compare state actors’ approaches to non-centralized media to those of regular users. This paper advances the study of non-centralized media in two key ways. First, it centers non-traditional users of media, like state agents and corporate actors. Second, it advances the use of leaked information in academic research as a method to reduce the problems of uncertainty and incentives to misrepresent information from official channels.

Introduction

Technology continues to integrate into everyday life, with AI entering workflows and social media continuing to gain importance as the locus of political and social debate and conversation. This integration between corporate owned technology and personal speech, creativity, and daily work life has raised concerns about surveillance and the amount of intimate, personal information that is continually being harvested for both profit and data to develop future technology. Many people have sought solace in free and open source software (FOSS) as a way to continue to enjoy the benefits of a technology-enabled lifestyle while avoiding the pitfalls of data harvesting, targeted advertising, and potential social and political repression. This trend has made its way into the world of non-centralized social media. For years, corporate owned platforms have dominated the media space; however, as companies have pivoted to a walled garden style of management, where features and development are focused on keeping users within the space, many people have sought alternative platforms. They leave walled gardens to escape the "[...] behavioral manipulation, pervasive surveillance, and capricious governance that characterizes large-scale centralized social platforms" (Kissane and Kazemi 2024). Unfortunately, while federated systems offer extraordinary opportunities to explore new forms of digital governance, their open nature makes them uniquely vulnerable to surveillance actors. A mixture of technologies enables both the rapid sharing and harvesting of information across the fediverse. The same standards, relays, and platform specific APIs that drive a thriving community also populate information for surveillance and malicious actors. This begs the important question: with all of the open and available information on the fediverse, who is likely to access it, and how will that information be used?

With Mastodon and the Fediverse entering a "killer hype cycle" after the purchase of Twitter (currently X), scholarship on the benefits of a non-centralized architecture for everyday end users has grown. However, a gap remains where other types of users are concerned. Following the framework set forth by Zulli, Liu, and Gehl (2020), I will expand on the implications of an open federated system and the state surveillance actors that often accompany them. First, I will offer a brief history of Mastodon and the technological imaginary that has developed alongside it. Then, the BlueLeaks dataset, a series of leaked internal emails, training, and notices from US fusion centers will be used to predict the response of state actors when approaching surveillance via the Fediverse. This article provides an analysis of fediverse users who have specialized use cases for social media, and have methods of engagement that differ from a traditional end user. As the Fediverse is accessible for any actor to use, host, and develop, understanding the influence that non-traditional users hold over the system as a whole is necessary to predict the future of non-centralized free media.

Mastodon and the socio-technical imaginaries of non-centralized media

Alternatives to corporate owned social media have been growing in popularity as the downsides of corporate owned social media have been making daily headlines. Most notably, Elon Musk's acquisition of Twitter, and subsequent renaming to X, pushed many users off of the platform. These users, angered by the ability of one person to gain complete control over what appeared to be a digital commons, sought out platforms arranged to prevent future attempts to centralize and monopolize participation. Many of these users turned to Mastodon, a popular non-centralized micro-blogging platform (Zia et al. 2023). This platform, built on a mixture of the ActivityPub protocol and its own APIs, allows users to join existing servers, or instances, to facilitate their browsing and content creation. Mastodon does not only exist as a simple alternative to sites such as Twitter, rather its unique mix of protocols, social practices, and ethics of participation make it a significantly different community than is found on corporate owned sites. Initially, Mastodon may seem quite empty to new users as content is not offered to them. The ActivityPub protocol that Mastodon uses for its server-to-server communication features an outbox for users to post content they wish for the fediverse to see. Individual users must follow one another for their server to seek out the content in others' outboxes and pull it into their feed (Ilik and Koster 2019). While the local content within an instance is populated automatically, it is often a new experience for users to need to seek out the content that they wish to see on their own feed. Thus, new users often run up against a form of culture shock when engaging with non-centralized social media for the first time.

User imaginaries

For fediverse proponents, Mastodon was built upon and maintains an ideal of a freer internet. It is a reaction against corporate entities which gain their power by housing data in silos and obscuring content engagement behind proprietary algorithms. Rather, the ActivityPub protocol that Mastodon uses for federation emphasizes the ability of users to create their own instances where they may control content moderation. Those who choose not to host can move between existing servers while maintaining their followers, lowering barriers to exit. This arrangement sets up its own set of unique trade-offs. Primavera De Fillipi notes that the benefits of a centralized system come at the cost of trust in the central authority. The central authority may encrypt or restrict access to data to ameliorate security concerns, but this means they have absolute access to that data (Filippi 2016). While users are aware their data is being harvested in a closed system, they are aware of the authority that is parceling and selling their data. Additionally, platforms may promote content for discoverability, but this comes at the cost of their ability to suppress or remove content.

On the other hand, non-centralized systems such as Mastodon allow for the freedom to own and remove content, at the cost of the content being widely visible (Filippi 2016). In this openness, Mastodon instances have developed their own community standards, creating groupings of communities that follow similar approaches to acceptable content, interaction, and moderation policies (R. W. Gehl and Zulli 2023, Rozenshtein 2022). In this case, trust can be placed in a far fewer number of hands, most often a minimal admin and moderator team. Often, instances will decide on these standards democratically, allowing instance users input into the content that they wish to populate their instance, increasing their trust in the community (Laux and Erd 2023). This bucks a trend of imposing the responsibility for content control on either individuals who are oftentimes powerless to actually modify the information on their feeds, or on to a mixture of outsourced labor and algorithms, which are far removed from the context of the interaction (Jang, Barrett, and McGregor 2024). These community developed standards often differ from the policies of corporate social media and can be tested as larger numbers of users begin to migrate away from corporate sites. In response to increased traffic, Mastodon moderators have been more vigilant for harassment on their servers, with some instances choosing to switch to an approval system rather than open enrollment in a bid to prioritize quality of interaction over quantity (Anaobi et al. 2023, Zulli, Liu, and R. Gehl 2020). This move emphasizes that, while Mastodon comes from a history of users seeking internet spaces free from the influence and control of corporations, this is not directly equivalent to a desire for an online space with no rules or moderation. While there do exist instances that prioritize unrestricted speech over content moderation, they tend to be blocked, or defederated, by instances that value the protection of their members. These factors make for a better intra-instance experience for users, but internal server governance is unable to shield users from the drawbacks of server-to-server communication.

There are several ways to access information on Mastodon. As previously mentioned, Mastodon currently implements the ActivityPub standard. This ensures that the data sent between servers is formatted in the common JSON format with uniform identifiers that any ActivityPub platform can ingest. Any entrant to the fediverse can be ensured that as long as they can read and write information according to the standard they will be able to share information across platforms. This standardization of information also means that any actor who wishes to scrape information has very clear instructions for accessing and understanding the information. Although the ActivityPub standard does have designations between information that is meant to be posted on a public timeline and that which is meant to be private to the user, it is the role of the instance to respect that. Even information marked with the private flag can be easily accessed if sent to an instance whose purpose is scraping information. Servers also bear the burden of user authentication. The World Wide Web Consortium, or W3C, the creators of the ActivityPub standard, marked authentication as outside of the scope of the standard (*ActivityPub* 2024). The organization links to community developed recommendations, but development and implementation of any form of authentication may change from instance to instance (*ActivityPub/Primer/Authentication Authorization - W3C Wiki* 2024). Additionally, the existence of relays serves as another layer of abstraction for siphoning data from the Fediverse. Many instances can subscribe to a relay by sharing their own information in exchange for pulling in content from a much larger swath of the fediverse. While relays do have an approval process, any newly created instance that wishes to surveil en mass would be quick to attempt to gain access.

Servers do have one line of defense in the form of defederation, or refusing to send or accept information from a specific instance. When an instance identifies a different server as malicious, it can block that server's access to its data. While defederation serves as a powerful tool for blocking malicious communities, the speed and inconspicuous nature of data harvesters is likely to render it ill-equipped as a method for protecting the privacy of data. This problem is likely to grow with the fediverse, as it may be difficult to distinguish the self-hosted server of a regular user who wants high-level control over their own experience from the self-hosted server of a malicious actor scraping federated data.

In addition to the implications of the use of the ActivityPub standard, Mastodon also implements its own internal application programming interface, or API, that allows clients to access information from instances without hosting an instance of their own. By making an API call to a specific instance, anyone could quickly access that instance's public timeline, public accounts, posts featuring specific hashtags, and more. The upside of the API call is that instances that respect the private tag on data will not send the information out in response. On the other hand, the barrier to entry to scraping data is far lower with the API. Recently, Maven, a social media startup, has come under fire for ingesting millions of Mastodon profiles and posts (R. Gehl 2024a). Although it remains contested whether they utilized the Mastodon API or their own server, this event demonstrates that small to mid-sized entities, or even individuals, can access large parts of the fediverse.

Pointing out the openness of data is not meant to be a critique of federated media. As it stands, the open sharing of information is necessary to maintain a web that many people can connect to on their own terms. This section does directly contradict one common pitch given to entice people to the fediverse, however; that a move out of a walled garden will make surveillance and data harvesting more difficult for corporations to do. It is clear that users with limited technical expertise could surveil the fediverse, but there remains the question of motivations for doing so. While surveillance by corporations is a major and pressing concern, this article focuses mainly on the surveillance tactics that state agents are likely to employ. In this endeavor, it is useful to take a step back from the underlying protocols that enable the fediverse to exist and instead explore the commonly held ideas of what can and does happen on the platforms.

State imaginaries

The commonly told vision of the fediverse as a collection of communities sustained and moderated by their members to ensure equitable access and participation lies in stark contrast of the version that is portrayed by state surveillance agents: a lawless space where anti-state violent activism can breed. The next section further develops the justification for painting a bleak picture of the states' concept of the fediverse, but at this point it is helpful to understand how preconceived notions of a system can direct responses to it. Jasanoff's concept of socio-technical imaginaries helps to untangle these opposing views of non-centralized social media. While users may have their experiences shaped by the underlying technologies that form the backbone for Mastodon, outside perceptions need not relate to the capabilities or uses of the technology. Most lawmakers and surveillance actors do not use federated media and thus may lack an understanding of the ways that the protocols shape their behavior. These officials perform their duties based on public discourse, aggregated reports, and hypothetical scenarios. Preconceived notions of what technology might be used for are often influential in molding legislation and policy making surrounding a technology (Jasanoff and Kim 2015). Outside of the fediverse, tech development often follows the ideals of neo-liberalism, where centralized technocrats create solutions that can disrupt and improve upon community developed solutions (Ferrari 2020). In the corporate realm, centralized authorities are empowered to solve complex issues, such as the appropriate standards for content moderation, and government actors have a clear authority to engage with. These imaginaries are baked into the institutions that negotiate the relationships between state agents and corporate media entities (Mager and Katzenbach 2021). For example, state agencies collect and distribute contact information for liaisons within social media companies. In these relations, there are clear entry points for state agents to request information and to serve legal documents. This entry point does not exist in a non-centralized system. State agents may or may not be able to contact instance administrators, and those same administrators may not be within their legal jurisdiction, requiring extensive coordination with other nation states to acquire or preserve information. Additionally, as many Mastodon instance administrators envision the fediverse as a free and open technology, they are far more likely to be opposed to government intrusion on the space (Ferrari 2019). This topology conflicts with expectations for social media and challenges conventional logic for state agents.

Non-centralized media as an open-source project also challenges traditional logics of online participation. Users may participate through the development of standards and platforms, in addition to building community through socialization. It was not uncommon for community developed features to be incorporated into social media applications. Bucking this trend, companies like Reddit and Twitter have been making the push to incorporate features of third-party apps, then cut off third-party access to their systems, forcing users onto their native applications (Peters 2023, Wiggers 2023). On the other hand, users of open-source projects can become known in their communities for developing features, in addition to providing content on the platform. Some projects, such as Scuttlebutt, even pair experienced developers with interested community members, helping them grow both their education in software development and their participation in the community (Mannell and Smith 2022). These systems can present a far less legible picture to law enforcement as there are too many entry points to hold individuals liable. Thus, state agents are likely to champion legal strategies that simplify responsibility for server actions, such as concentrating liability onto hardware owners.

The open web works because it is open. The plurality of places and ways to enter and engage with the fediverse at the will of the user is part of what drives the virtuous communities which thrive on federated platforms. That same openness allows for mass data harvesting and surveillance. That being the case, members of the fediverse will show concern for the methods and frequency that surveillance is employed on them and their information. In the following section, the BlueLeaks dataset is used to understand the vision that state surveillance agents have of the fediverse. This dataset addresses several key implications for surveillance. First, it provides insight into the actions of both national agencies and local police departments. Although federal agencies are likely to be better equipped to scrape data en masse, the most pervasive surveillance can occur from local forces targeting a fediverse member in their community. Second, by analyzing training data, this dataset provides a view into the methods that build and sustain a technological imaginary that can persist in the state surveillance apparatus. Finally, the leaked nature of the data overrides the incentives of state agents to misrepresent or downplay the extent of their surveillance into the fediverse.

BlueLeaks and Fusion Centers

In 2020, 270 gigabytes of data were leaked from U.S. Intelligence Fusion Centers. This information contained training materials developed by the fusion centers, lists of sign-ups for training from local police forces, inventory lists for certain centers, lists of email requests for center services from locals and local police forces, alerts for danger to officers, suspicious persons reports, reports on potential protests, and more[1]. This dataset offers a rare look at the internal workings of the United States intelligence community and provides a chance to understand how state forces engage, conceptualize, and train others about decentralized media.

Fusion centers were founded post 2001 with the idea that enhanced communication between intelligence agencies and local law enforcement would improve their ability to prevent future acts of terror. Fusion centers were conceived as a node for the aggregation of data from all federal and local agencies. The nodes could then disseminate the combined information from larger sectors of the intelligence community. Additionally, these centers allow national agencies to harness local resources and coverage, trading training and broad intelligence for local intelligence and monitoring capabilities. Currently, there are over 70 recognized fusion centers in the United States covering all states and territories (*Fusion Center Locations and Contact Information — Homeland Security* 2024). While there have been reports that local officers have found the training and notices on officer safety and potential hazards key to performing their jobs safely (Lewandowski and Jeremy G Carter 2017), the success of these centers has been contested. Researchers have found that increased surveillance does not inherently lead to better security outcomes (Degli Esposti, Ball, and Dibb 2021). Additionally, some local police forces have refused to rely on the fusion center network in favor of leveraging their own surveillance tactics (McQuade 2019). In addition to pushback from academics and local forces, fusion centers have been criticized by national lawmakers. In 2012, senators Levin and Coburn led a congressional

[1] The full database is hosted by Distributed Denial of Secrets https://ddosecrets.com/wiki/BlueLeaks. For the less technically inclined, Micah Lee of DDOSecrets has authored a book detailing the process of downloading and navigating this dataset, Hacks Leaks and Revelations, without which this paper would not be possible

inquiry into the effectiveness of fusion centers, summarized well by the Committee on Homeland Security and Governmental Affairs activity report (*112th Congress Committee on Homeland Security and Governmental Affairs* 2024):

> Federal funding provided by the Department of Homeland Security (DHS) to State and local intelligence "fusion centers" had not yielded significant useful information to support Federal counterterrorism efforts. Among other problems, the Coburn-Levin report showed that the fusion centers produced intelligence that was of uneven quality, was often untimely, and sometimes endangered civil liberties, and showed that DHS did not effectively monitor the use of Federal funds provided to State and local fusion centers, which sometimes made questionable expenditures. In addition, the report determined that senior DHS officials were aware of the problems hampering effective counterterrorism work with the fusion centers, but did not always inform Congress of the issues, nor ensure the problems were fixed in a timely manner.

Despite criticism from congressional members, academics, police officers, and citizens, fusion centers continued their growth, receiving over 400 million dollars in funding in 2021 (*2021 National Network of Fusion Centers Assessment: Summary of Findings* n.d.). In spite of this public investiture, available information about fusion centers and their activities can be sparse and self-contradictory. While some may see intelligence centers as justifiably more secretive, there appears to be widespread confusion about basic facts, like their role in the intelligence community and the actual scope of their activities (Jeremy G. Carter, Lewandowski, and May 2016). Studies show that this lack of transparency hinders potential local acceptance of their actions by citizens who support national security objectives if goals and methods are clear (Westerlund, Isabelle, and Leminen 2021). Fusion centers were unable to maintain this lack of transparency, however, as in 2020 a hack now known as BlueLeaks exposed a wide variety of information about fusion center activities, members, and opinions.

On June 6th, 2020, a hacktivist from the group Anonymous was able to gain access to 251 law enforcement websites, all of which shared vulnerabilities by virtue of being developed by the same company, Netsentinal (Lee 2020a). This leak gave the hacker back-end access to all files hosted by the website, including information that, although declassified, was sectioned off for access only by law enforcement and corporate surveillance partners. Of particular interest to journalists and social movement organizations were the Suspicious Activity Reports, or SARs. These reports are put out daily by some fusion centers and compile a list of potential activities for local law enforcement to monitor. In 2020, protests led by the Black Lives Matter organization were heavily tracked, with the Northern California Regional Intelligence Center (NCRIC) sending out a full list of potential protests twice daily (Lee 2020b). SAR reports are often paired with the social media posts used to verify the information, demonstrating the importance of Open Source Intelligence (OSINT) to the monitoring activities of the fusion centers.

This revelation justified the concerns of many people, who fear that the information they put out on social media is being collected and analyzed for use in political repression. While the move away from corporate media may seem an obvious solution to this problem, this data also revealed that information flows from everyday users back to surveillance agencies. SARs attach the source of their information, many of which were by everyday citizens acting as "surveillance deputies", sending information to the fusion center as suspicious (Brayne, Lageson, and Levy 2023). While instance admins and moderators may pursue strategies to eliminate surveillance deputies, such as applications to join instances, it is unlikely that this strategy could fully eliminate end-user driven reporting. Users often make compromises between resisting and encouraging surveillance, depending on their views in a particular situation. Users know they operate within a complex ecosystem of surveillance and are constantly renegotiating their role within that sphere (Lyon 2017).

The potential for any user to report information to authorities has been shown to produce chilling effects on participation and user engagement. While many are aware that social media is a highly public space, it is still perceived to be a sort of private public, where information is not expected to be monitored, harvested, or sold (Nissenbaum 2009). When users become aware that they are being actively surveilled, their participation in the community drops off (Scott 2017). Users who leave corporate social media to join non-centralized media are likely more concerned about surveillance and data harvesting. With the mass scraping and sale of data being the preeminent threat to privacy, more personalized and targeted surveillance by community members may present a surprise. Users migrating to non-centralized media will likely be forced to address the possibility of both state actors joining communities under false pretenses, and also community members acting as surveillance deputies when they feel a line has been crossed.

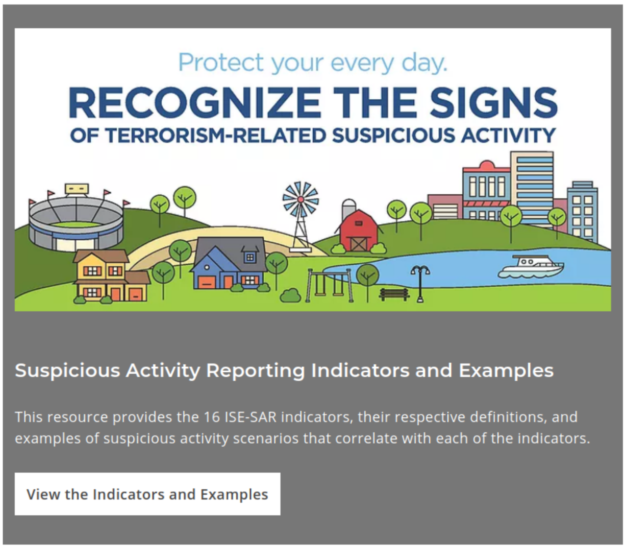
Protect your every day.
RECOGNIZE THE SIGNS
OF TERRORISM-RELATED SUSPICIOUS ACTIVITY

Suspicious Activity Reporting Indicators and Examples

This resource provides the 16 ISE-SAR indicators, their respective definitions, and examples of suspicious activity scenarios that correlate with each of the indicators.

View the Indicators and Examples

Figure 1: "See something say something" style messaging from the Oklahoma Fusion Center webpage. Accessed 06/09/2024

State actors recognize the perception of privacy that non-centralized media brings. For them it represents a security risk: actors may be able to plan and organize acts of terror utilizing the benefits of connecting to like-minded people over the internet and the ability to send messages and files with secure encryption. While there has been some evidence supporting the idea that online spaces can act as pathways to radicalization that then leads to cases of violence and murder (Mann et al. 2023, Karell et al. 2023), the deterministic nature of this discourse has been contested (Colley and Moore 2022). Without clear evidence, state agents still appear to be concerned about decentralized instances that have previously been harnessed by those already holding extreme views. Most notably, the extreme right-wing group Gab forked Mastodon in a bid to gain access to the fediverse. This was seen as a move to circumvent restrictions from corporate owned play stores, aiming to keep the site accessible to a wider range of users. This move directed federal attention to federated social media, as the users they were already invested in tracking had now moved to the platform. The September 5-11, 2019 issue of the Counterterrorism Digest makes clear the concerns that state actors have about decentralized media[2]:

> In July 2019 developers at USPERGab forked-when developers copy source code from one software package and create a distinct and separate piece of software-the source code of Mastodon, a decentralized social network. By forking Mastodon, Gab has become a decentralized social platform and is subsequently more resistant to take-down efforts. In addition, Gab is now also available on mobile phones through their own app store.
>
> Terrorist and violent extremists are increasingly investigating the use of decentralized networks to facilitate the spread of terrorist and hateful material online. Only recently, so-called ISIS supporters have experimented with Mastodon and Jihadoscope recently discovered ISIS supporters experimenting with Mastodon (see inset). This further underlines that terrorists continue to utilize open-source software available to anyone.
>
> ISIS has been long keen on finding a viable alternative to Telegram and in theory, they could follow the same path as Gab has done in building a take-down resistant platform by forking open source-code(s). Increasingly terrorists and violent extremists are building their own software applications and much of this depends on re-using existing code and software that was originally developed under the open-source model and published for everyone to re-use and modify.

At this pivotal moment, surveillance actors are beginning to develop strategies for engaging with federated social media at large. Zulli, Liu, and Gehl have developed a framework for understanding how federated systems change interactions between users. They identified Topology, Abstraction, and Scale as general categories that illustrate how the technical aspects of federation shape behavior. In the following sections, I apply this framework to surveillance actors to illustrate how their unique position is altered by federation as well.

Topology

As may be expected, the most notable difference in network topology between federated and centralized systems is the distributed nature of nodes. While all traffic for centralized corporate social media has a touch point on their servers, the same does not remain true for

[2] This issue can be found at: BlueLeaks/bostonbric/files/DDF/(U) 11 September 2019 NCTC CT Digest.pdf

federated social media. Although there are different standards in different federated environments, information exchange among servers is the most common method of communication. For ActivityPub, one of the most popular standards, each user has an inbox and outbox that are facilitated by their instance. Content they wish to push to the fediverse is placed in their outbox, and their inbox pulls in information from the outboxes of users they have followed. In order for an inbox to successfully pull in information from a user on a different instance, there must be a point of connection between their instances. This is why widely connected servers allow users a wide range of choices, and isolated servers may seem desolate. In a robustly connected instance, as portions of the network drop off traffic can be rerouted through any federated server to ensure access to information. This also underscores the importance of instances remaining in good standing with one another. If an instance is cut off from well connected nodes, they may lose access to significant portions of the fediverse.

For end users, this means that the network encourages interest-based sociality instead of attention-seeking behavior, as economic incentives for the production of content have fallen away in favor of social benefits (Zulli, Liu, and R. Gehl 2020). Federated media often has no way of algorithmically serving content to users. Content is pulled in chronologically from followed accounts, so users must instead focus on building network connections with others who share their interests to ensure interest-based content reaches them. Rather than solely focusing on content that is likely to meet algorithmic criteria, and thus be served to more people, users must build networks through direct interaction. This lack of a central authority serving content to users provides a potential benefit to state surveillance actors. Many point to algorithmically served content as a funnel for users to learn more about fringe groups, and to slowly be exposed to more of their messaging. If a user consistently engages with content, most algorithms will continuously serve it to them, allowing that content to cover more and more of their feed. By working within the factors for content discovery, organizations can ensure their content is shown to more and more people. If anti-state or violent organizations of interest to state actors are able to reach wider audiences, they may be able to increase their membership or inspire individuals to take violent action on their behalf.

One major benefit of corporate social media's algorithmic model is the ability to go viral. If posts or accounts gain steam, there always remains the possibility that content could be served to millions of viewers in a very short period of time, at little to no cost to an organization. This is part of what makes corporate social media so attractive to organizations. For state agents, the ability to go viral and reach a potentially impressionable audience poses a security risk. In this respect, federated systems are likely to act in favor of state security actors over the interests of organizers.

Network topology cuts both ways for state agents. Without a centralized location to monitor, the time, data, and surveillance necessary to monitor the online activities of entities of interest increase dramatically. This results in the view that federated social media can become a breeding ground for violent and anti-state speech, where discovering and monitoring dissent is difficult. Much like the dark web, which has many legitimate political and social uses (R.

Gehl 2018), all activities on federated serves run the risk of being seen as illicit. Additionally, although many Mastodon instances engage in a shared moral code of ethics (R. W. Gehl and Zulli 2023), there remains the possibility that servers that do allow or encourage hate speech continue to exist in the greater network that comprises the fediverse. Even though these actors may be blocked, or defederated, by many or most instances on the fediverse it can be difficult for state agents to trace their connections and see what other instances or accounts are engaging with them. Although obtaining information from the fediverse is quite easy, as previously discussed, it can be much more difficult to synthesize that information into actionable intelligence. This task becomes much simpler in centralized systems, where a company can fully investigate the posts and connections of a profile of interest.

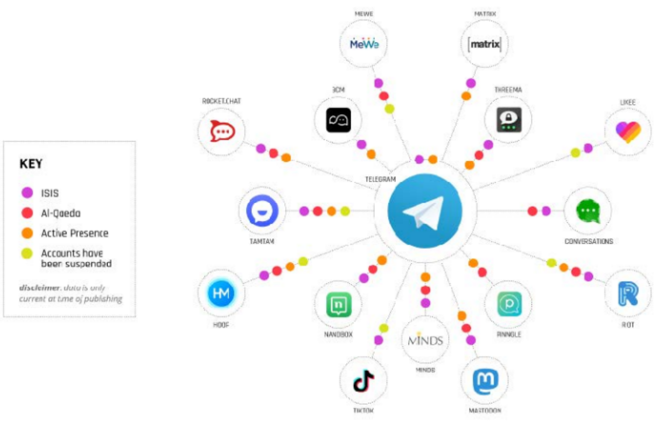

Terrorist social media graph **(FLASHPOINT)**

Figure 2: Info-graphic from the January 16-22 2020 issue of Counterterrorism Digest. It depicts Mastodon as a platform that has actively hosted ISIS and Al-Qaeda accounts. This diagram depicts a view of decentralization across platforms, rather than a focus on decentralization within platforms.

Additionally, without a central authority to denounce malicious actions and to take responsibility for removing it from the space, there could remain perception among outsiders that illegal actions are accepted by the fediverse. By focusing on the servers that host illicit and illegal content, state actors receive a shaded view of the totality of interactions occurring on federated systems. A more nuanced understanding of interactions between federated servers would be necessary for surveillance actors to determine the connectedness of any single instance in the larger fediverse. As the Counterterrorism Digest graphic in figure 2 shows, surveillance actors are less interested in understanding decentralization within platforms, but rather look at organizations first, then take an interest in all platforms that they spread to. This means that any platform (or in the case of the fediverse, grouping of platforms that share a method for interconnecting) can become suspect.

Without an understanding of the topology of Mastodon and other services sharing information via standardized protocols, it is likely that intelligence agencies will attempt to monitor many networks of servers without first exploring the extent to which malicious instances are federated. This is likely to lead to increased surveillance of everyday users, harming both their rights to privacy and their experience on Mastodon (Landwehr, Borning, and Wulf 2023), in addition to wasting funds and time for state surveillance actors.

Abstraction

The framework utilized in this paper deals with abstraction in the software sense; systems become more abstract the further that users get from machine code. Mastodon and ActivityPub software tends to get users far closer to their actual data than corporate owned systems, who make their income by parceling out and and reselling it to third parties. While state actors may be closer to public data than they are in corporate owned systems, they are further from direct access to servers. Data flagged as private is not federated through relays and is sent directly from server to server. Therefore, to harvest private data state agencies must either get a person of interest to send that data to a server that they are hosting, or obtain access to their home server or the home server of people they communicate with. Because data is distributed among many smaller servers, it can be more difficult for state agents to locate and secure legal access to servers of interest. This brings state agents further from the information they seek, as they no longer have a single point of contact, such as a corporate office, nor do they have predefined relationships of compliance to requests for information. This section further expands on the use of corporate contacts for the surveillance apparatus of the United States, and the ways that federated communication complicates this relationship.

For the average end user, social media operating on the ActivityPub protocol is far more legible than corporate social media. As Zulli et al. point out, the sites are developed by users, have no algorithmically served content, and often have open moderation policies that are explained to users before they are harshly enforced (Zulli, Liu, and R. Gehl 2020). Users see the content that they have subscribed to in chronological order and are often able to have discussions with moderators if their posts are being removed from the instance. This eliminates any confusion about the reasons information is being served to, or removed from them. For state agents, federated systems may prove to be far more abstract than corporate owned systems. State agents rely heavily on their corporate partners to collect, preserve, and present information. Unlike an end user, who may struggle to connect with a corporation to find additional information on blocked content and suspended profiles, state agents have designated contacts at social media companies to ensure their requests are recorded and responded to in a timely manner. It can be difficult to put an exact number on these requests. Proton, primarily an email service dedicated to privacy, posts the number of requests for user data that it complies with yearly. This number has hovered around 6,000 - 7,000 for the last three years, although Proton points out that all of these requests come through the Swiss government, and it is unclear if they are to support investigations for other nations (*Proton Transparency report* 2024). Statista reports that in the first half of 2023, Google fielded over 81,000 requests from the United States government alone (*Statista* 2024). While this report does not touch on the number of these requests that Google complied with, it shows that federal agencies commonly reach out to companies that process large amounts of data.

Many fusion centers, in their role as educators for local police forces, collate corporate contact information and provide templates for officers to use, ensuring that requests to preserve and serve information have the proper language to prompt compliance. Although the responsibilities for content moderation are currently in dispute[3], social media companies generally understand their legal requirements when dealing with requests from law enforcement. These interactions are structured, and fusion centers work to ensure agencies can interact with social media companies using standardized templates.

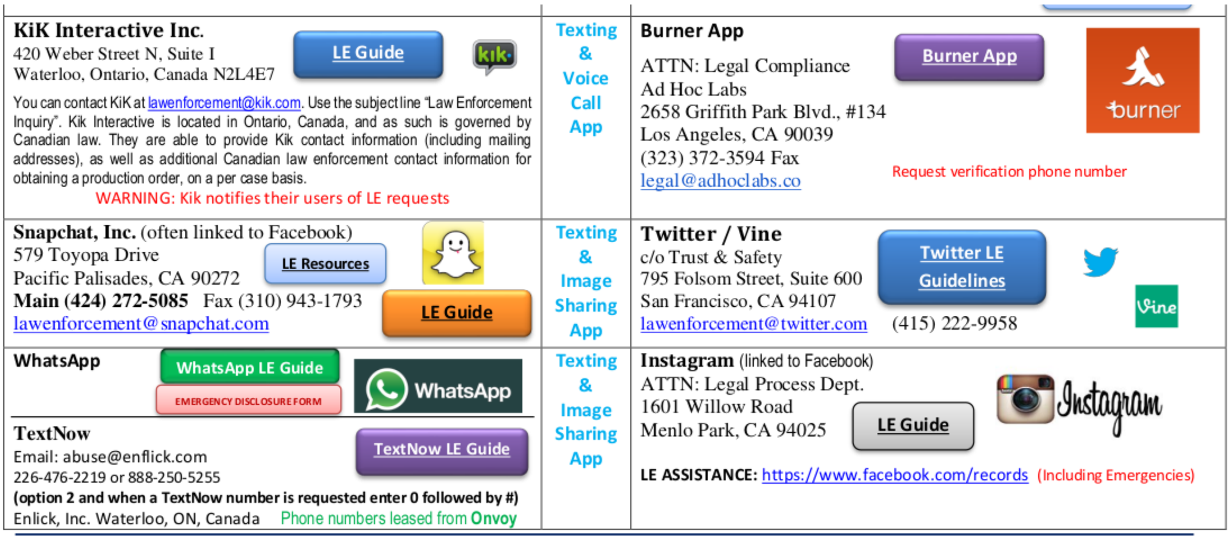

KiK Interactive Inc. 420 Weber Street N, Suite I Waterloo, Ontario, Canada N2L4E7 LE Guide You can contact KiK at lawenforcement@kik.com. Use the subject line "Law Enforcement Inquiry". Kik Interactive is located in Ontario, Canada, and as such is governed by Canadian law. They are able to provide Kik contact information (including mailing addresses), as well as additional Canadian law enforcement contact information for obtaining a production order, on a per case basis. WARNING: Kik notifies their users of LE requests	Texting & Voice Call App	**Burner App** ATTN: Legal Compliance Ad Hoc Labs 2658 Griffith Park Blvd., #134 Los Angeles, CA 90039 (323) 372-3594 Fax legal@adhoclabs.co	Burner App Request verification phone number
Snapchat, Inc. (often linked to Facebook) 579 Toyopa Drive Pacific Palisades, CA 90272 **Main (424) 272-5085** Fax (310) 943-1793 lawenforcement@snapchat.com LE Resources LE Guide	Texting & Image Sharing App	**Twitter / Vine** c/o Trust & Safety 795 Folsom Street, Suite 600 San Francisco, CA 94107 lawenforcement@twitter.com	Twitter LE Guidelines (415) 222-9958
WhatsApp WhatsApp LE Guide EMERGENCY DISCLOSURE FORM **TextNow** Email: abuse@enflick.com 226-476-2219 or 888-250-5255 **(option 2 and when a TextNow number is requested enter 0 followed by #)** Enlick, Inc. Waterloo, ON, Canada Phone numbers leased from **Onvoy** TextNow LE Guide	Texting & Image Sharing App	**Instagram** (linked to Facebook) ATTN: Legal Process Dept. 1601 Willow Road Menlo Park, CA 94025 **LE ASSISTANCE:** https://www.facebook.com/records (Including Emergencies)	LE Guide

Figure 3: An excerpt from an Alabama Fusion Center guideline collecting and publishing the information department contacts for many common social media apps and phone carriers. This guide links to pre-drafted request templates for the preservation of information. Found at BlueLeaks/alabamafusioncenter/files/DDF/LE Technology Investigations Resource Guide REVISED 01-15-18CF R.PDF

Although instance moderators may have the same visibility to law enforcement and state agents as they do to end users, their likelihood to comply with orders to preserve and export information is far less clear. Instance administrators may be subject to different nation's laws or may be more willing to seek legal recourse to delay back-end access to their server. As there is conflict in international regulation surrounding media, this makes enforcement of local laws far more complicated for simple surveillance (Gstrein 2020). Coordinating with international agencies to seize servers increases expense, time, and logistics to requests that may be simple when working with a corporation based within the agency's own nation. This is not to say that corporations are unable, or in few select circumstances unwilling, to seek remedy for police requests in court. However, the existence of liaisons for law enforcement points to an amicable relationship between the two.

While it may be theoretically possible for intelligence agencies to collect and analyze social media information on their own, it appears that they currently rely heavily on their corporate partners to assist them in this endeavor. As the September 28th through October 4th 2017

[3] Recent supreme court cases NetChoice, LLC v Moody and NetChoice, LLC v Paxton concern disputes over content moderation affecting First Amendment rights. As of the writing of this paper, both cases have been sent back to the lower courts to reconsider (*Moody v. NetChoice, LLC* 2024, *NetChoice, LLC v. Paxton* 2024)

issue of the Counterterrorism Digest, a weekly compilation of open-source information with commentary issued by the U.S. National Counterterrorism Center, states[4]:

> Outreach and partnerships with private-sector social media/technology companies may facilitate early recognition of and coordination on suspicious activities for further vetting, potentially disrupting radicalization, recruitment, and attack plotting. Established partnerships, which can also help develop and ensure familiarity of "tripwires" for reporting suspicious material/content, is an example of how the private sector can assist authorities in thwarting terrorist activities.

It appears that intelligence agencies across the board encourage corporate partnerships to assist with sifting through massive amounts of collected data. As it is unknown whether instance administrators will comply with requests without being compelled by law, one potential tactic for state actors to continue mass surveillance would be to scrape information through platform specific APIs, personal instances, or impersonation on existing servers. This would require additional expenditure in either staff and storage or in allocations to a third party to scrape and analyze the data. While it is both possible and plausible that one to all of those options will be taken, they still represent a considerable hurdle for state actors that does not exist with their corporate partners. Additionally, this leaves state surveillance actors increasingly dependent on potentially unresponsive or inaccurate models and algorithms (Eldridge, Hobbs, and Moran 2018). Although it is likely that law enforcement will be able to access federated servers through a combination of intimidation and legal force, this is not a substitute for willing corporate partners. Much like fusion centers partner with local police forces to expand the flow of information up the surveillance pipeline, corporate partners are also necessary to identify potential threat sources and alert national surveillance centers.

A second point of abstraction for law enforcement and surveillance agencies is content moderation. As previously mentioned, the burden of moderation is a hotly contested issue. In the U.S., social media companies have protections under both the Digital Millennium Copyright Act and Communications Decency Act, which grant them safe harbor to publish content without being legally liable for its contents. This section allows companies to remove content, without requiring them to do so. In the U.S., there are currently debates over First Amendment rights of corporations and whether moderation policies constitute free speech, which may disrupt this arrangement. This is not contained to the United States, however. Recently Canada has also been experimenting with legislation to impose stricter requirements for social media companies to respond to flagged harmful content. As Robert Gehl points out, these laws are far more complicated for non-centralized media where instances may live in different nations than their administrators, and there is little to no income to support the technology and processes demanded by regulations (R. Gehl 2024b). While these regulations may greatly impede the growth of federated social media, they may also increase the red tape for intelligence agencies

[4] This issue can be found in the following file path in the BlueLeaks dataset: BlueLeaks/bostonbric/files/DDF/(U)CTDigest4October910.pdf

that wish to interact with them. As the fediverse grows, so to does the potential for unclear legal situations, making surveillance increasingly difficult to conduct in a legal manner.

Scale

For normal users of Mastodon, Zulli, Liu, and Gehl found that Mastodon scales on the number of instances rather than the number of users. They point out that user counts are very important to the business model of corporate social media companies, whereas Mastodon instance administrators were more focused on the quality of engagement on their servers. This led to administrators considering the possibility of restricting the number of users if the quality of interaction began to diminish. The authors point out that this leads to the importance of horizontal growth, increasing the connections between a large number of servers.

Theoretically, a fully connected fediverse would not pose as many issues for state surveillance actors as their agents or surveillance deputies would be able to connect to many other servers without the need to establish a profile on every single one of them. However, Mastodon instances do not actually prefer to federate across all other instances. Rather, communities develop where instances tend to federate with others that share their moderation policies and goals. Therefore, instances that allow hate speech and illegal activity are likely to be blocked, or defederated, by instances that value the protection of their members. This results in communities developing along a model of "conventional federalism", where they agree to federate with instances with similar moral codes of conduct and agree as a federation to block violating instances (R. W. Gehl and Zulli 2023). Even as the scale of the fediverse grows, users may not actually be exposed to more harmful content if conventional federalism remains strong. This feeds into a dilemma of state forces posed by Xu: state forces prefer that actors are visible to be apprehended, but prefer they are undiscoverable so that their messages remain unseen and cannot grow dissidence (Xu 2020). This problem worsens for certain non-centralized platforms, such as blockchain based technologies. As the blockchain is thought to be immutable, state forces cannot scrub dissident information found on these protocols (Mott 2019).

Large federal agencies with more employees and funding may be able to offset the costs of the collection and analysis of data. Services using ActivityPub are often open systems, with all nodes accessible by default. Timelines are published and accessible, and agencies with the storage and technical experience necessary can easily index vast swaths of the fediverse. However, fusion centers were founded due to the failings of centralized large data analysis. The Colorado Intelligence Center claims that "The mission of the Colorado Information Analysis Center is to serve as the focal point within the state for receiving, analyzing and sharing threat-related information among private sector, local, tribal, and federal partners."[5] These centers are founded on the idea that local information is of importance to the intelligence community. Being able to ingest and index the fediverse en masse turns surveillance on all of its members, while potentially yielding no actionable intelligence for the agencies involved. Due to the unclear links between pluralities of servers, following the trend of using local sources to monitor media

[5] Located at BlueLeaks/ciacco/files/DDF0000/607.pdf

sites and push the information upstream may be more difficult in federated systems.

The connections between instances are of the utmost importance for the flow of information to Mastodon users. Users will only be able to see content if their network includes the user or instance creating it. This results in a hodgepodge of overlapping information highlighted by Christina Dunbar-Hester in her comically named article, "Showing your ass on Mastodon"(Dunbar-Hester 2024). Dunbar-Hester coins the term "lossy-distribution" for the way that information fails to bridge the gap between servers. One particularly striking example is of users who were not able to see all of the replies to their own posts, and thus missed a portion of the conversation developing on a thread that they created. While this would be unthinkable in a centralized corporate social media, it appears that it is a common, although often unnoticed, occurrence on Mastodon. While this is impactful for everyday users (particularly those with an interest in donkeys in Dunbar-Hester's case), it is a significant risk in the eyes of intelligence agencies.

Another complication for surveillance agencies is that interaction within the fediverse occurs across many different services. While this paper uses Mastodon as the most common service using ActivityPub, many services allow their users to share information via the protocol, and that number is growing. Most recently, popular content publication software providers Flipboard and Ghost have enabled cross-posting content to the fediverse (Perez 2024, Flipboard 2024). In addition to many services plugging into the fediverse, they may not all support the same feature set. Although services must conform to a certain amount of the ActivityPub standard to share information, there are often add-on or platform specific aspects that cannot be shared. For example, Mastodon has a "boost" feature, similar to a Re-Tweet, that many non-Mastodon federated services cannot handle. This problem is complex enough to confound many fediverse developers and would represent a significant hurdle to surveillance actors attempting to gather information through the fediverse. These state actors must not only be aware of a presence on the fediverse, but also be aware of the nuances of the protocol and ensure they are utilizing compatible services. As each additional service adds federation support, thus increasing the scale of activity through the fediverse at large, they further complicate the ability of a state surveillance agency to gain a comprehensive view of federated activity.

Surveillance agencies rely on a complex community of citizens, local authorities, and corporations to collect and analyze intelligence (Martin, Van Brakel, and Bernhard 2009, Lewandowski and Jeremy G Carter 2017). This community relies on users closer to the source of information to be able to identify, document, and then pass the information up to authorities. These users may be agents of local police forces or partner companies, but are often citizens who encounter information in their daily lives that concerns them. This same strategy may not apply to federated systems. Users on the most connected instances are unlikely to encounter dangerous information, as these instances follow conventions that block harmful content and defederate from instances that allow it. Still, it can be difficult to locate instances that specifically host harmful information, as they are likely to have closed membership and will not be widely shared in the rest of the fediverse. While the solution to this issue may be to discover the existence of instances on corporate social media, where information is shared more widely, and then to bring surveillance directly to that instance, it seems likely that the ease of scraping information will result in mass surveillance of everyday users.

As of 2020, when the dataset for this paper was obtained, it is not clear whether U.S. intelligence agencies took a nuanced stance on Mastodon or the fediverse. An Open Source Intelligence Bulletin from the Central Florida Intelligence Exchange [6] accurately describes Mastodon as a series of servers sharing information on a standard protocol, but goes on to claim that this makes it ideal for violent extremists and anarchist extremists. Although this was a short write-up for law enforcement, it shows that the depth of information reaching local surveillance deputies lacks nuance about the connections between servers, or complications that horizontal scaling presents for surveillance. This lack of understanding of the social aspects that follow the growing scale of the fediverse points to potential privacy violations for users, and ineffectual monitoring for state agencies.

Conclusion

Social media has a profound effect on the way that individuals interact with their friends and community. While much of this is funneled through a small number of platforms, interest in more open, non-corporate systems has grown. Mastodon has surged in popularity, exposing far more people to the concept of federated systems. As the future of these systems seems bright, scholars must understand the implications of federated media for state agents and vice versa. It is important for fediverse proponents to understand that the interests of these parties are not monolithic with respect to decentralization. State agents lose willing corporate partners that collect and synthesize data for them. On the other hand, they lose the looming threat of algorithmically driven disinformation campaigns that seek to incite violence. Communication around federated systems should seek to inform both users and governing bodies about the efforts of content moderation that make individual instances and federated covenants safe and hospitable places for online interaction. Additionally, promoters of federated systems using the current ActivityPub standard should be more circumspect in promoting the fediverse as safe from surveillance or privacy violations.

There remain future versions of the decentralized web that sacrifice some of the ease of access that ActivityPub provides for better protections against data harvesting and a decreased need for trust in a specific server. Projects on the horizon promote decentralized identities as an evolution for the fediverse, where individuals are able to control their online identity, and more easily switch between providers. These systems currently offer much more autonomy at the level of the individual user, at the cost of being more technically difficult to set up and maintain. Whether or not these systems become as popular as server-based arrangements, addressing user concerns over data protection will be a key issue for the future of federated systems.

[6] This issue can be found in the following file path in the BlueLeaks dataset: BlueLeaks/seffc/files/DDF/CFIX OSINT Potential Use of ‘Mastodon Social’ as Twitter Alternative 18-12-122.pdf

As this paper shows, surveillance is already occurring on the fediverse. In addition to understanding the implications of topology, abstraction, and scale for users and state agents, it is imperative for scholars to understand the perceptions of these technologies. Policies and approaches to technology are often based on the imaginaries surrounding them, rather than their realized uses. This paper advances the use of leaked data for analyzing the perceptions of state agents in relation to emergent technology. In this case, leaked data protects against incentives to misrepresent the scale and direction of surveillance. This data also reveals more closely the operating views of state surveillance actors, as it includes the training data disseminated to local agencies. Understanding the concept of decentralization held by state surveillance agencies is necessary to predict and uncover surveillance practices that may be hidden from public view.

Finally, this paper paves the way for further work expanding on the non-homogeneous views of state agents in relation to non-centralized mass media. This paper has focused on the training material and support provided by central agencies and fusion centers to local authorities. Local agencies are likely to form their own views of the uses of non-centralized media, and conduct surveillance and monitoring accordingly. Additional sources are necessary to create a more comprehensive view of the socio-technical imaginaries that will be perpetuated by the policies developed in the coming years. For example, the BlueLeaks dataset also contains requests for resources, funding, and surveillance submitted by local authorities. This information is likely to contrast in part with the information found in training and teaching material, allowing researchers insight into contesting views on the role of non-centralized media in the surveillance community. To fully understand the extent that surveillance will affect the fediverse, researchers must examine all facets of social and political discourse that may influence state agent perspectives.

References

112th Congress Committee on Homeland Security and Governmental Affairs (2024). URL: https://www.hsgac.senate.gov/wp-content/uploads/Activities_Report_-_112th_Congress_PSI_Excerpt_1.pdf (visited on 06/09/2024).

2021 National Network of Fusion Centers Assessment: Summary of Findings (n.d.). Tech. rep.

ActivityPub (2024). URL: https://www.w3.org/TR/2018/REC-activitypub-20180123/#security-considerations (visited on 09/01/2024).

ActivityPub/Primer/Authentication Authorization - W3C Wiki (2024). URL: https://www.w3.org/wiki/ActivityPub/Primer/Authentication_Authorization (visited on 09/01/2024).

Anaobi, Ishaku Hassan et al. (Apr. 2023). “Will Admins Cope? Decentralized Moderation in the Fediverse”. In: *Proceedings of the ACM Web Conference 2023*. arXiv:2302.05915 [cs], pp. 3109–3120. DOI: 10.1145/3543507.3583487. URL: http://arxiv.org/abs/2302.05915 (visited on 04/09/2024).

Brayne, Sarah, Sarah Lageson, and Karen Levy (Dec. 2023). *Surveillance Deputies: When Ordinary People Surveil for the State*. en. SSRN Scholarly Paper. Rochester, NY. URL: https://papers.ssrn.com/abstract=4661112 (visited on 06/09/2024).

Carter, Jeremy G., Carla Lewandowski, and Gabrielle A. May (Sept. 2016). “Disparity Between Fusion Center Web Content and Self-Reported Activity”. en. In: *Criminal Justice Review* 41.3, pp. 335–351. ISSN: 0734-0168, 1556-3839. DOI: 10.1177/0734016816651925. URL: http://journals.sagepub.com/doi/10.1177/0734016816651925 (visited on 04/05/2024).

Colley, Thomas and Martin Moore (Jan. 2022). “The challenges of studying 4chan and the Alt-Right: ‘Come on in the water’s fine’”. en. In: *New Media & Society* 24.1. Publisher: SAGE Publications, pp. 5–30. ISSN: 1461-4448. DOI: 10.1177/1461444820948803. URL: https://doi.org/10.1177/1461444820948803 (visited on 04/30/2022).

Degli Esposti, Sara, Kirstie Ball, and Sally Dibb (Sept. 2021). “What’s In It For Us? Benevolence, National Security, and Digital Surveillance”. en. In: *Public Administration Review* 81.5, pp. 862–873. ISSN: 0033-3352, 1540-6210. DOI: 10.1111/puar.13362. URL: https://onlinelibrary.wiley.com/doi/10.1111/puar.13362 (visited on 04/12/2024).

Dunbar-Hester, Christina (Mar. 2024). “Showing your ass on Mastodon: Lossy distribution, hashtag activism, and public scrutiny on federated, feral social media”. en. In: *First Monday*. ISSN: 1396-0466. DOI: 20240313092153000. URL: https://firstmonday.org/ojs/index.php/fm/article/view/13367 (visited on 03/14/2024).

Eldridge, Christopher, Christopher Hobbs, and Matthew Moran (Apr. 2018). “Fusing algorithms and analysts: open-source intelligence in the age of ‘Big Data’”. en. In: *Intelligence and National Security* 33.3, pp. 391–406. ISSN: 0268-4527, 1743-9019. DOI: 10.1080/02684527.2017.1406677. URL: https://www.tandfonline.com/doi/full/10.1080/02684527.2017.1406677 (visited on 04/05/2024).

Ferrari, Elisabetta (Jan. 2019). “‘Free country, free internet’: the symbolic power of technology in the Hungarian internet tax protests”. en. In: *Media, Culture & Society* 41.1, pp. 70–85.

ISSN: 0163-4437, 1460-3675. DOI: 10.1177/0163443718799394. URL: http://journals.sagepub.com/doi/10.1177/0163443718799394 (visited on 04/18/2024).

Ferrari, Elisabetta (Apr. 2020). "Technocracy Meets Populism: The Dominant Technological Imaginary of Silicon Valley". en. In: *Communication, Culture and Critique* 13.1, pp. 121–124. ISSN: 1753-9129, 1753-9137. DOI: 10.1093/ccc/tcz051. URL: https://academic.oup.com/ccc/article/13/1/121/5739591 (visited on 04/18/2024).

Filippi, Primavera De (2016). "The interplay between decentralization and privacy: the case of blockchain technologies". en. In: *Journal of Peer Production* 7.

Flipboard (Aug. 2024). *Now People on Flipboard Can Follow Anyone in the Fediverse*. en-US. URL: https://about.flipboard.com/fediverse/follow-anyone-in-the-fediverse/ (visited on 09/15/2024).

Fusion Center Locations and Contact Information — Homeland Security (2024). en. URL: https://www.dhs.gov/fusion-center-locations-and-contact-information (visited on 06/09/2024).

Gehl, Robert (2018). *Weaving the Dark Web*. Cambridge, Mass.: The MIT Press.

— (June 2024a). *Maven Ain't So Mavenly*. en. URL: https://fossacademic.tech/2024/06/12/Maven.html (visited on 07/07/2024).

— (Apr. 2024b). *Reading the Online Harms Act with my Fediverse Admin Hat On*. en. URL: https://fossacademic.tech/2024/04/10/Online-Harms-Act.html (visited on 07/07/2024).

Gehl, Robert W. and Diana Zulli (Dec. 2023). "The digital covenant: non-centralized platform governance on the mastodon social network". en. In: *Information, Communication & Society* 26.16, pp. 3275–3291. ISSN: 1369-118X, 1468-4462. DOI: 10.1080/1369118X.2022.2147400. URL: https://www.tandfonline.com/doi/full/10.1080/1369118X.2022.2147400 (visited on 04/09/2024).

Statista (2024). *Google: data requests from federal agencies H1 2023*. en. URL: https://www.statista.com/statistics/273501/global-data-requests-from-google-by-federal-agencies-and-governments/ (visited on 09/15/2024).

Gstrein, Oskar Josef (2020). "Mapping power and jurisdiction on the internet through the lens of government-led surveillance". en. In: *Internet Policy Review*. ISSN: 21976775. DOI: 10.14763/2020.3.1497. URL: https://policyreview.info/articles/analysis/mapping-power-and-jurisdiction-internet-through-lens-government-led-surveillance (visited on 04/11/2024).

Ilik, Violeta and Lukas Koster (June 2019). "InformationSharing Pipeline". en. In: *The Serials Librarian* 76.1-4, pp. 55–65. ISSN: 0361-526X, 1541-1095. DOI: 10.1080/0361526X.2019.1583045. URL: https://www.tandfonline.com/doi/full/10.1080/0361526X.2019.1583045 (visited on 04/09/2024).

Jang, Heesoo, Bridget Barrett, and Shannon C. McGregor (Apr. 2024). "Social media policy in two dimensions: understanding the role of anti-establishment beliefs and political ideology in Americans' attribution of responsibility regarding online content". In: *Information, Commu-*

nication & Society 27.6. Publisher: Routledge _eprint: https://doi.org/10.1080/1369118X.2023.2234970, pp. 1047–1072. ISSN: 1369-118X. DOI: 10.1080/1369118X.2023.2234970. URL: https://doi.org/10.1080/1369118X.2023.2234970 (visited on 07/06/2024).

Jasanoff, Sheila and Sang-Hyun Kim (2015). *Dreamscapes of Modernity: Sociotechnical Imaginaries and the Fabrication of Power*. en. University of Chicago Press. ISBN: 978-0-226-27652-6 978-0-226-27649-6 978-0-226-27666-3. DOI: 10.7208/chicago/9780226276663.001.0001. URL: https://www.bibliovault.org/BV.landing.epl?ISBN=9780226276663 (visited on 07/06/2024).

Karell, Daniel et al. (Mar. 2023). ""Born for a Storm": Hard-Right Social Media and Civil Unrest". en. In: *American Sociological Review*. Publisher: SAGE Publications Inc, p. 00031224231156190. ISSN: 0003-1224. DOI: 10.1177/00031224231156190. URL: https://doi.org/10.1177/00031224231156190 (visited on 03/21/2023).

Kissane, Erin and Darius Kazemi (Aug. 2024). *Findings Report: Governance on Fediverse Microblogging Servers*. en-US. Tech. rep. Publication Title: Manubot. Manubot. URL: https://fediverse-governance.github.io/ (visited on 10/25/2024).

Landwehr, Marvin, Alan Borning, and Volker Wulf (Jan. 2023). "Problems with surveillance capitalism and possible alternatives for IT infrastructure". en. In: *Information, Communication & Society* 26.1, pp. 70–85. ISSN: 1369-118X, 1468-4462. DOI: 10.1080/1369118X.2021.2014548. URL: https://www.tandfonline.com/doi/full/10.1080/1369118X.2021.2014548 (visited on 04/04/2024).

Laux, Lea and Laszlo Erd (2023). "Trust as the Elephant in the Room – Security Evaluation of Decentralized Online Social Networks with Mastodon". en. In.

Lee, Micah (July 2020a). *Hack of 251 Law Enforcement Websites Exposes Personal Data of 700,000 Cops*. en-US. URL: https://theintercept.com/2020/07/15/blueleaks-anonymous-ddos-law-enforcement-hack/ (visited on 06/09/2024).

— (Aug. 2020b). *How NCRIC Keeps Tabs on Black Lives Matter Protesters*. en-US. URL: https://theintercept.com/2020/08/17/blueleaks-california-ncric-black-lives-matter-protesters/ (visited on 06/09/2024).

Lewandowski, Carla and Jeremy G Carter (May 2017). "End-user perceptions of intelligence dissemination from a state fusion center". en. In: *Security Journal* 30.2, pp. 467–486. ISSN: 0955-1662, 1743-4645. DOI: 10.1057/sj.2014.38. URL: http://link.springer.com/10.1057/sj.2014.38 (visited on 04/05/2024).

Lyon, David (2017). "Surveillance Culture: Engagement, Exposure, and Ethics in Digital Modernity". en. In.

Mager, Astrid and Christian Katzenbach (Feb. 2021). "Future imaginaries in the making and governing of digital technology: Multiple, contested, commodified". en. In: *New Media & Society* 23.2, pp. 223–236. ISSN: 1461-4448, 1461-7315. DOI: 10.1177/1461444820929321. URL: http://journals.sagepub.com/doi/10.1177/1461444820929321 (visited on 04/22/2024).

Mann, Marcus et al. (Jan. 2023). “Unsorted Significance: Examining Potential Pathways to Extreme Political Beliefs and Communities on Reddit”. en. In: *Socius: Sociological Research for a Dynamic World* 9, p. 237802312311748. ISSN: 2378-0231, 2378-0231. DOI: 10.1177/23780231231174823. URL: http://journals.sagepub.com/doi/10.1177/23780231231174823 (visited on 10/02/2024).

Mannell, Kate and Eden T. Smith (July 2022). “Alternative Social Media and the Complexities of a More Participatory Culture: A View From Scuttlebutt”. en. In: *Social Media + Society* 8.3. Publisher: SAGE Publications Ltd, p. 20563051221122448. ISSN: 2056-3051. DOI: 10.1177/20563051221122448. URL: https://doi.org/10.1177/20563051221122448 (visited on 09/07/2022).

Martin, Aaron K, Rosamunde E Van Brakel, and Daniel J Bernhard (Apr. 2009). “Understanding resistance to digital surveillance: Towards a multi-disciplinary, multi-actor framework”. en. In: *Surveillance & Society* 6.3, pp. 213–232. ISSN: 1477-7487. DOI: 10.24908/ss.v6i3.3282. URL: https://ojs.library.queensu.ca/index.php/surveillance-and-society/article/view/3282 (visited on 04/11/2024).

McQuade, Brendan (2019). *Pacifying the Homeland: Intelligence Fusion and Mass Supervision.* Oakland, CA: University of California Press.

Moody v. NetChoice, LLC (2024). en-US. URL: https://www.scotusblog.com/case-files/cases/moody-v-netchoice-llc/ (visited on 07/07/2024).

Mott, Gareth (Feb. 2019). “A Storm on the Horizon? “Twister” and the Implications of the Blockchain and Peer-to-Peer Social Networks for Online Violent Extremism”. In: *Studies in Conflict & Terrorism* 42.1-2, pp. 206–227. ISSN: 1057-610X. DOI: 10.1080/1057610X.2018.1513986. URL: https://doi.org/10.1080/1057610X.2018.1513986 (visited on 08/30/2019).

NetChoice, LLC v. Paxton (2024). en-US. URL: https://www.scotusblog.com/case-files/cases/netchoice-llc-v-paxton/ (visited on 07/07/2024).

Nissenbaum, Helen (2009). *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford: Stanford University Press. ISBN: 978-0-8047-5236-7.

Perez, Sarah (2024). *Substack rival Ghost confirms it will join the fediverse in 2024 — TechCrunch.* URL: https://techcrunch.com/2024/04/22/substack-rival-ghost-confirms-it-will-join-the-fediverse-in-2024/ (visited on 09/15/2024).

Peters, Jay (June 2023). “Reddit - Dive into anything”. In: *The Verge*. URL: https://embed.reddit.com/r/reddit/comments/145bram/addressing_the_community_about_changes_to_our_api/?embed=true&ref_source=embed&ref=share&utm_medium=widgets&utm_source=embedv2&utm_term=23&utm_name=post_embed&embed_host_url=https%3A%2F%2Fwww.theverge.com%2F2023%2F6%2F9%2F23755640%2Freddit-api-changes-apps-apollo-shut-down-ama-spez-steve-huffman (visited on 07/06/2024).

Proton Transparency report (2024). en. URL: https://proton.me/legal/transparency (visited on 09/15/2024).

Rozenshtein, Alan Z. (Nov. 2022). *Moderating the Fediverse: Content Moderation on Distributed Social Media.* en. SSRN Scholarly Paper. Rochester, NY. URL: https://papers.ssrn.com/abstract=4213674 (visited on 11/23/2022).

Scott, Jeramie D (2017). "Social Media and Government Surveillance: The Case for Better Privacy Protections for Our Newest Public Space". en. In: *Journal of Business and Technology Law* 12.2.

Westerlund, Mika, Diane A. Isabelle, and Seppo Leminen (Apr. 2021). "Perspectives from Higher Education: Applied Sciences University Teachers on the Digitalization of the Bioeconomy : The Acceptance of Digital Surveillance in an Age of Big Data". en. In: *Technology Innovation Management Review* 11.3, pp. 32–44. ISSN: 19270321. DOI: 10.22215/timreview/1427. URL: https://timreview.ca/article/1427 (visited on 04/11/2024).

Wiggers, Kyle (Jan. 2023). "Twitter officially bans third-party clients after cutting off prominent devs". en-US. In: *TechCrunch.* URL: https://techcrunch.com/2023/01/19/twitter-officially-bans-third-party-clients-after-cutting-off-prominent-devs/ (visited on 07/06/2024).

Xu, Xu (2020). "To Repress or to Co-opt? Authoritarian Control in the Age of Digital Surveillance". In: *American Journal of Political Science* 65.2, pp. 309–325. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/ajps.12514 (visited on 07/06/2024).

Zia, Haris Bin et al. (Feb. 2023). *Flocking to Mastodon: Tracking the Great Twitter Migration.* arXiv:2302.14294 [cs]. DOI: 10.48550/arXiv.2302.14294. URL: http://arxiv.org/abs/2302.14294 (visited on 06/06/2023).

Zulli, Diana, Miao Liu, and Robert Gehl (July 2020). "Rethinking the "social" in "social media": Insights into topology, abstraction, and scale on the Mastodon social network:" en. In: *New Media & Society.* Publisher: SAGE PublicationsSage UK: London, England. DOI: 10.1177/1461444820912533. URL: https://journals-sagepub-com.mutex.gmu.edu/doi/metrics/10.1177/1461444820912533 (visited on 07/23/2020).