M^{3}AAWG | Messaging Malware Mobile
Anti-Abuse Working Group

May 27, 2016

Via Electronic Submission to http://apps.fcc.gov/ecfs/

Federal Communications Commission
FCC Wireline Competition Bureau, Competition Policy Division
445 12th Street SW
Washington, DC 20554

RE: WC Docket No. 16-106
U.S. FCC Protecting the Privacy of Customers of Broadband and Other Telecommunications Services

Background:

On April 1st, 2016, the U.S. Federal Communications Commission (FCC) released a notice of proposed rule making (NPRM) on the issue of "Protecting the Privacy of Customers of Broadband and Other Telecommunications Services." This proposal laid out several new definitions for what is considered private information for a telecommunications customer and a number of new rules that internet service providers (ISPs) within the United States would have to abide when dealing with that information.

This response reviews the proposal with regard to its impact on the anti-abuse mission that is the core of the Messaging, Malware and Mobile Anti-Abuse Working Group (M^{3}AAWG).

Executive Summary

There is a school of thought that says privacy and security are incompatible and that in order to have one, you have to intrinsically give up the other. This is a philosophy that we at M^{3}AAWG explicitly reject. We have long been driven by the philosophy that protecting our users from abuse requires we protect BOTH their privacy and their security. We recognize such a "win-win" scenario is not easy and that it requires careful design and appropriate safeguards to ensure it will not be abused. Indeed many of our member companies take steps today to protect consumers from a myriad of online threats and abuse which also serve to enhance consumer privacy.

To quote the M^{3}AAWG website: "The Messaging, Malware and Mobile Anti-Abuse Working Group (M^{3}AAWG) is where the industry comes together to work against bots, malware, spam, viruses, DoS attacks and other online exploitation." As such, we offer a vetted community where anti-abuse practitioners can gather to discuss current issues, talk about what works (and what does not), and document best practices to share with the larger internet community.

All of this work is based on the fundamental idea that only by sharing information with each other about the threats we encounter can we, as a community, hope to fight off the myriad of threats and challenges we face. Indeed, this sort of data sharing and collaboration is a bedrock of the anti-abuse community as a whole, going back to some of the very first anti-abuse solutions on the internet.

With this as background, M^{3}AAWG is concerned that the NPRM, as it is currently written, does not provide sufficient safeguards and carve-outs to allow U.S. ISPs to continue to provide the critical collaboration and data sharing that we all need to fight the good fight. We recognize that the NPRM proposes to exempt from the privacy regime ISP uses or disclosures of customer information that "protect the rights or property of the provider, or to protect users and other providers from fraudulent, abusive, or unlawful use of, or subscription to, broadband services".[1]

While we appreciate this language, we are concerned that it is too narrow and vague, and could therefore potentially interfere with a considerable amount of anti-abuse work. This concern is heightend by the extraordinary breadth of the data elements that the FCC proposes to include in the definition of "customer proprietary information" (CPI) subject to the rules, including IP addresses, MAC IDs and domain information. The data elements proposed to be covered by the rules are central to our work, even though they do not inherently or automatically identify any specific person and we have no interest in any such identity. Our comments are aimed at ensuring that the Commission takes into consideration circumstances in which the proposed rules could inhibit our anti-abuse efforts.

To this end, we will lay out a number of examples of work we do today that could be significantly curtailed, if not out-and-out banned, under the NPRM as written. We also identify use cases where the NPRM as written appears to be in contradiction with other existing legislation. We hope that by enumerating these use cases, we can bring greater visibility to the anti-abuse work that is done today and encourage the commission to draft language that protects this work going forward.

DNS Blackhole Lists (DNSBLs)

A DNS Blackhole List is a list of IP addresses and blocks that have demonstrated a history of sending spam email. Please see this Wikipedia article for more background (https://en.wikipedia.org/wiki/DNSBL). A variety of organizations and groups build and maintain DNSBLs. Some are licensed for profit, others are operated on a non-profit basis, and still others are maintained internally by large email receivers. Many date back to the origins of internet anti-abuse work from the late 1990s.

What all DNSBLs hold in common, however, is that they gather information from email receivers, ISPs, and other resources to identify the sources of spam messages. Much of the data that is collated together to form a DNSBL is sourced from ISPs, both within and without the U.S. If ISP use and disclosure of suspect domains or suspect IP addresses for DNSBL purposes becomes subject to the proposed privacy rules, the degree these lists can offer network providers timely and useful information could be severely compromised. It should also be noted that in compiling these lists, there are invariably instances of "false positives" – i.e., domains or IP addresses that may initially appear to be sources of abusive activity but upon further scrutiny are not (or are determined to have been spoofed in some manner). In most cases, these "false positives" are rooted out before they are included in a DNSBL. However, since "false positive" information ultimately turns out not to be necessary to "protect" networks and users against spam and other abusive conduct, ISPs may

[1] *Protecting the Privacy of Customers of Broadband and Other Telecommunications Services,* Notice of Proposed Rulemaking, WC Docket No. 16-106, FCC 16-39, ¶ 115 (rel. April 1, 2016) (http://transition.fcc.gov/Daily_Releases/Daily_Business/2016/db0401/FCC-16-39A1.pdf).

be concerned that they would risk incurring penalties for any such disclosures. As a result, the volume and frequency of the abuse-related information they share may diminish significantly, to the detriment of the public.

A good example of data sourced from ISPs is the Spamhaus PBL, or "Policy Block List." The PBL attempts to identify all IP space assigned to dynamic customers, such as those who subscribe to most common residential broadband services. Additional information is found here (https://www.spamhaus.org/pbl/). The PBL is a tremendously valuable resource that helps email providers identify incoming mail that may contain spam, phishing or malware.

The PBL is a great example of a service that ISPs provide for the common good. ISPs get no additional benefit from helping keep the PBL accurate, but others benefit tremendously by having up-to-date and correct information from ISPs. Based on the linking language in the NPRM, the proposed rules might discourage participation or substantially reduce the quantity and quality of information available.

Feedback Loops

Feedback loops are a system used by most email providers to get information back to the email sender about messages sent from their network. They allow senders to better manage the email coming out of their network and reduce the amount of malicious spam, phishing and malware. A good description of how feedback loops work is available here: https://en.wikipedia.org/wiki/Feedback_loop_(email).

Feedback loops operate by sending complaints about traffic from a sender's IP blocks back to that sender. Customer complaints in the form of clicking on a "this is spam" button, or a like relay, are sent back to the sender. This gives senders timely alerts on bad traffic emanating from their system, allowing them to quickly identify problems and resolve them. Such loops quickly identify and stop the malicious use of otherwise legitimate mail platforms for sending bad traffic.

Unfortunately, based on the language in the NPRM, the feedback loop operator (be it the ISP itself or someone contracted by the ISP) could potentially be required to get the customer's approval before relaying any messages back via a feedback loop. If approval is required for every message, it would effectively render the loop useless. If a single approval is required, this would still dramatically lower the efficacy of such a loop, as many messages would not be sent due to a lack of approval. Again, we recognize there may be an intent that the exception for CPI uses aimed at protecting against "abusive" conduct should apply here. But we are concerned that the risk of penalties may discourage ISPs from feeding information into the loop – or seeking assistance from a third-party – absent definitive certainty that the email address or IP blocks being referenced are associated with malicious traffic.

ARIN Registration

All IP blocks allocated by ARIN are required, as part of the ARIN contract, to have ownership information entered into the shared WHOIS project (http://whois.arin.net/ui/). WHOIS allows anyone on the internet to lookup which ISP is assigned a given IP block and has been a fundamental building block of how IP addresses are assigned since the internet was first created. It is required as part of the contract between the ISP and ARIN when obtaining the IP block in the first place.

When many ISPs register an IP block, they provide additional information in order to assist others who might run into an issue emanating from the block. Examples are blocks that are explicitly labeled for dynamic customers or labeled for a specific region. For instance, the 24.24.0.0/14 IP block assigned to Road Runner (former TWC) states in the comments that this block is used by customers in Ohio or the Carolinas (https://whois.arin.net/rest/net/NET-24-24-0-0-1/pft?s=).

There is no concrete link to prevent the abusive use of service in connection with publishing and labelling IP blocks in ARIN, so it is not apparent this exception (or any other in the NPRM) applies. This is simply standard practice and has been for some time. So the question becomes is this sort of activity a potential violation of the CPI protections and/or linkability standards discussed in the NPRM? If so, does that mean an ISP will have to potentially violate the terms of their contract with ARIN in order to not violate the NPRM? Which takes precedent?

Reverse DNS

As part of the specification for DNS, every IPv4 address in use on the internet should have a DNS PTR record associated with it (often referred to as a reverse DNS entry). Most ISPs include a small amount of information in this reverse DNS entry to help classify the IP address. Based on the CPI protection for domain information and/or the linkability language in the NPRM, there are concerns that this information in the PTR records would need to be changed, which could have a negative impact on overall security. Here again though, it may be difficult to tie such a change to a specific and imminent security threat, which then raises questions about whether the protection of property exception could be applied here.

Examples:

96.241.229.24 is an IP address assigned to Verizon. The reverse DNS for this is "pool-96.241-229-24.washdc.fios.verizon.net", which tells you it is a dynamic customer located in the Washington D.C. metro area connected by Fiber.

24.0.1.1 is an IP address assigned to Comcast. The reverse DNS for this is "c-24-0-1-1.hsd1.nj.comcast.net", which tells you that this is a dynamic customer located in New Jersey.

Academics and Researchers

As with any area of cutting edge innovation, much of the new thinking and research that leads to technological advances for anti-abuse starts with security researchers. Sometimes these researchers are associated with an academic institution, and other times they are independent workers, deriving income from their research activities. In all these cases, it is common, and in fact encouraged, for ISPs and service providers to partner with these researchers in order to provide them access to the aggregate data they need to validate their research.

There are many examples of successful security techniques and businesses spawned through the work of these researchers. In most cases, research derives a possible new technique. The researchers then develop the technique and a solution to utilize it, then need to test it on live data. Sometimes this takes place through industry coalitions, such as M^{3}AAWG. Other times, researchers partner with individual ISPs, or a small set of ISPs, to access the real world data required to validate their findings. Researchers often repeat this last step several times with different ISPs, using different iterations of the solution, as a way to take a rough idea and translate it into a working product.

At its heart, the security goal of any such work is to take a large amount of data in aggregate and identify the key information in that data that allows for identification of a problem, infection or other security concern. As a result, researchers attempt to use anonymous data to identify signs of a security problem, either on the host ISP network or pointing at signs on another network.

All of this would be at considerable risk, given the significantly higher barriers to data sharing and data access the new NPRM would create. There is a threshold question of whether this work could even take place outside the context of a specific threat to the network or malicious attack.

Beyond that basic question, the NPRM does not allow any data to be categorized as "aggregate" – and therefore outside the consent requirement – if it is "reasonably linkable" to a device. Even though the data that M^{3}AAWG works with could rarely, if ever, be linked to a specific person almost all of this data may in some fashion be considered "reasonably linkable" to a device because it all ultimately derives from devices. So the data flows from ISPs that we depend upon for anti-abuse research could be adversely affected.

In addition, the NPRM spells out quite clearly that any entity to which the ISP may furnish aggregate data would be contractually obligated not to try and re-identify from this larger data set. But the researchers and academics that we work with may attempt to discern patterns or commonalities in data sets in order to help identify strategies, tactics and defensive measures – and those patterns and commonalities may include sifting for device categories or characteristics associated with a particular threat vector or attack scenario. Would that be considered an attempt at re-identification that we would be contractually bound to forego? The concepts of device linkability and re-identification in the NPRM are broad and ISPs may be unwilling to open themselves up to this risk. As such, ISPs would be dramatically less inclined to support such research and researchers due to the inherit risk this would now entail.

Takedowns

An unfortunate aspect of the always-connected nature of U.S. broadband customers is the attractiveness they provide to criminals looking to abuse their network connections and computational power. Such bots and botnets represent a massive and growing threat to the online security landscape, powering network abuses from the sending of spam and phishing to proxying illegal activity to generating denial of service attacks.

From time to time, an entity (sometimes law enforcement, sometimes industry) identifies a bot and a mechanism for disabling that botnet. The entity takes the appropriate legal steps to instruct the various ISPs to assist with the takedown of that botnet through some technical means, such as blocking network access to the botnet. In addition, the entity often asks that infected users be redirected (if possible) to an agreed upon list of tools in order to try and remove the bot infection. This last step is critical, because with most bot takedowns, the criminals behind the botnet will not only work actively to thwart the takedown attempt, they also may, after a period of time, attempt to re-establish their control over their bot empire.

These actions have been some of the more successful coordinated efforts by the online community to fight criminal activity. We recognize that the exception for CPI uses to protect against "abusive or unlawful" use of broadband service probably would accommodate notifying customers infected by an ongoing malicious attack – though the Commission should confirm this.

As written, however, the NPRM could require additional notification and approval steps to ensure that all potentially infected users do the final clean-up steps listed above – particularly in a scenario where the imminence of the botnet's threat to the network has passed, the potential for its re-establishment has not, and some affected end users require additional reminders. While ensuring the hygiene of potentially infected devices can have a significant negative effect on deterring further attacks, the potential need to obtain customer approval to provide such notifications and resources once an imminent threat has passed could lower the number of machines successfully mitigated, and thereby lower the success of the event.

Conclusion

As the above use cases have shown, many of the techniques and tools utilized today to fight online messaging abuse are predicated on the successful sharing of data elements between ISPs and other internet services – be it other ISPs or third-parties – that are categorized as CPI in the NPRM. These data exchange models work because they allow security professionals to share data with minimal friction. The NPRM as it is currently written would add considerable friction to these tools and mechanisms, preventing the exchange of key information, and therefore would significantly impair the efficacy of these existing tools. And even if the NPRM is specifically drafted to allow these current use cases, it will still prevent the creation of new security techniques and mechanisms that leverage data sharing models not currently envisioned.

Many of the problems identified here would be mostly (or perhaps fully) negated if the NPRM made clear that data elements identified as CPI – such as IP addresses and domain information – can be used without permission in circumstances where they do not identify any specific person because that is how the vast bulk of the information covered by these examples is used today. At a minimum, M^{3}AAWG recommends the NPRM language be modified so that information sharing to facilitate online safety and security be more explicitly and broadly carved out and authorized; there not be any limitations placed on sharing in this regard; and there not be a requirement that any such uses be tied to a specific or imminent network threat or malicious attack. Further, provisions incorporated into the NPRM should not limit the ISP's ability to gather and store data for the purposes of fighting online abuse.

About the Messaging, Malware and Mobile Anti-Abuse Working Group (M^{3}AAWG)

The Messaging, Malware and Mobile Anti-Abuse Working Group (M^{3}AAWG) is where the industry comes together to work against bots, malware, spam, viruses, denial-of-service attacks and other online exploitation. M^{3}AAWG (www.m3aawg.org) members represent more than one billion mailboxes from some of the largest network operators worldwide. It leverages the depth and experience of its global membership to tackle abuse on existing networks and new emerging services through technology, collaboration and public policy. It also works to educate global policy makers on the technical and operational issues related to online abuse and messaging. Headquartered in San Francisco, California, M^{3}AAWG is driven by market needs and supported by major network operators and messaging providers.

Thank you for the opportunity to submit these comments. We will be glad to respond to any questions. Please address any inquiries about our work to me, M^{3}AAWG Executive Director Jerry Upton, at jerry.upton@m3aawg.org.

Sincerely,
Jerry Upton, M^{3}AAWG Executive Director
Jerry.Upton@m3aawg.org