To: Matt Larson, VP of Research, ICANN Office of the CTO (OCTO)
From: Odeline MacDonald, ICANN Legal
Date: 16 February 2024

Re: Visible Interruption (VI) and Visible Interruption and Notification (VIN) - Privacy and review

Background information on VI and VIN

The Name Collision Analysis Project (NCAP) Discussion Group (DG) is nearing completion of its Study 2, which proposes multiple techniques to collect data and evaluate the risk of name collisions for new gTLD strings. In contrast, the 2012 new gTLD round used only one technique, Controlled Interruption (CI), for this purpose.

Two of the techniques proposed in Study 2 are “Visible Interruption” (VI) and “Visible Interruption and Notification” (VIN). Both techniques are similar, though VIN entails an extra step that provides the “notification” in the technique’s name.

For both techniques, a new gTLD string is delegated in DNS to a zone containing only wildcard A (IPv4) and AAAA (IPv6) records. These records direct any client who queries any name in the gTLD zone to a “Sinkhole Server” that rejects all protocols (in the case of VI) or a “Sinkhole and Response Server” that rejects all protocols except HTTP connections (in the case of VIN). With VIN, the Sinkhole and Response Server responds over HTTP with a message explaining the name collision situation. This explanatory message is the notification component referenced in the technique’s name. For example, if VIN were enabled for the TLD string “new-gtld” and a device attempted to connect to the URL “http://www.some-domain.new-gtld”, the recursive resolver on the device would resolve the name “www.some-domain.new-gtld” and the result would be the IP address of the Sinkhole and Response Server. The device would then connect to that server over HTTP and the server would return an explanatory page with information about name collisions.

In the lead up to the 2012 new gTLD round, the ICANN Security and Stability Advisory Committee (SSAC) examined alternative notification approaches, including one called a "honeypot" (an information security term of art). The SSAC, however, did not reach consensus in recommending the honeypot approach against the Controlled Interruption approach that was ultimately adopted by the Board.

Several participants in the NCAP DG, including those from OCTO, have raised concerns about the privacy implications of VI/VIN. With these techniques, traffic leaves a client device and reaches the sinkhole server, potentially divulging sensitive information. For example, HTTP requests can include parameters including user names and passwords and all protocol connection attempts to the sinkhole server reveal the IP address of the connecting device, which could possibly be a device being used directly by a human (as opposed to, for example, a server in a data center). Such activity may entail the collection and

processing of "personal data," which refers to any information relating, directly or indirectly, to an identified or identifiable natural person as defined under the General Data Protection Regulation (GDPR) and various other global regulations. In contrast, the IP address used for Controlled Interruption is a “loopback” address, which causes any packets sent to that address to return, or “loop back”, to the sending device itself without ever leaving the device.
It is worth noting that VI features a notification component akin to Controlled Interruption (CI), where applications cease to function. However, the mechanism differs from simply receiving a non-existent name answer from the root servers. Instead, applications are provided with an IP address for connection. In the case of VI, connection attempts are promptly rejected, triggering application-specific error messages. Similarly, in VIN, all connections except HTTP encounter rejection, resulting in application-specific errors.

In contrast, with CI, traffic remains confined within the client's device: the client application endeavors to connect to the CI address, a 'loopback' address that redirects all packets back to the client. Both VI and VIN reroute the application's connection attempt to the sinkhole server, thus exposing the attempt to anyone along the path, including the sinkhole server itself.

The purpose of this high-level privacy review is to assess the privacy and legal implications of VI and VIN as techniques for name collision data collection and end user notification.

Summary of the review
VI and VIN are techniques utilized to gather data for assessing the risk of domain name collisions. In the case of VIN, they also serve to notify end-users. Both VI and VIN, equipped with a notification feature akin to CI, have the drawback of causing application disruptions. Unlike CI, which maintains traffic within the client's device, VI and VIN reroute connection attempts to a sinkhole server, potentially exposing them to interception. While VI outright rejects connections, leading to application-specific errors, VIN only permits HTTP connections, resulting in similar errors. This stands in contrast to CI, where traffic remains within the client's device and is redirected back to it."

While VI/VIN can be useful tools for preventing conflicts, they also pose privacy and data protection related risks that should be considered. Such risks encompass the possible absence of a legal basis, as it might not be possible to rely on consent, legitimate interest or any other statutory justification as grounds for processing, depending on the type of personal data included in an HTTP request. Indeed, when weighing the interests of ICANN or other parties deriving benefits from conducting VI/VIN against the interests or fundamental rights and freedoms of data subjects concerned, ICANN’s or the third parties’ interests might be outweighed considering the significant risks and negative impacts on privacy and data protection, and potentially other fundamental rights and freedoms of data subjects.

Furthermore, the lack of transparency in data processing, the absence of data minimization practices, and the risk of exposing sensitive data of data subjects without legal basis make VI/VIN a concern regarding compliance with data protection laws. Apart from data protection, there are also risks related to disclosure of confidential information.

Entities conducting VI/VIN could mitigate these risks by implementing appropriate steps such as identifying a suitable legal basis for the processing of personal data, applying data minimization measures, ensuring transparency about their practices (including, clearly formulated notices), implementing appropriate security measures and conducting a data protection impact assessment, as might be required under applicable data protection laws such as the GDPR. Generally, compliance with the relevant data protection laws will be key to mitigate privacy and liability risks.

However, due to the unpredictable nature of the data collected and the potentially high residual risks associated with data collection and further processing, implementing these privacy safeguards appears to be challenging, if not impossible. This situation may lead to potential legal and reputational consequences for the entities conducting VI/VIN.

Privacy and data protection assessment

VI/VIN can raise privacy and data protection concerns as well as risks related to disclosure of confidential information. As it is not predictable which type of data will be collected and further processed, there are potential residual high risks associated with the data processing in this context. As it will be difficult or even impossible to implement suitable mitigation measures this could have legal and reputational consequences for the entities conducting VI/VIN.

1. **Identifying a legal basis for the processing of personal data**

The legal basis for processing personal data in the context of VI/VIN will depend on the specific circumstances and the applicable data protection laws. However, there are a few potential legal bases that may apply in this context:

- *Legitimate interests*: the most likely legal basis for processing personal data in the context of VI/VIN is legitimate interests. Legitimate interests may be invoked when the processing is necessary for the purposes of the legitimate interests pursued by the controller or a third party. An example of legitimate interest of the controller may be to assess the risk of a collision in order to prevent harm to individuals and property. However, this interest must respect a number of requirements:
    - it must be lawful, clear, real and present;
    - the processing must be a necessary for pursuing this interest (in other words no milder means should be available to achieve the desired result);
    - the interests of the data controller or a third party in processing personal data for VI/VIN must outweigh the interests or fundamental rights and freedoms of the data subjects, which means that a balancing test needs to be conducted. This balancing test would involve weighing the benefits of the processing against the potential risks for and negative effects on the data subjects and be conducted once, determining that in all such cases, the interest in notifying users about the collision outweighs their privacy interests.

-> See more information on the balancing test in the paragraph on safeguards.

- *Consent:* Another potential legal basis for processing personal data in the context of VI/VIN is consent. If the data processing involves collecting, using or sharing personal data for a specific purpose, and there is no other legal basis that applies, the data controller may need to obtain the data subject's explicit and informed consent before processing the data.
  -> This solution cannot be realistically implemented.

- *Legal obligation*: A third potential legal basis for processing personal data in the context of VI/VIN is compliance with a legal obligation. This may apply if there is a legal requirement for the data controller to process personal data for the purposes of assessing the risk of a collision.
  -> According to our knowledge, there is currently no legal obligation requiring using VI/VIN.

Legitimate interest as legal basis and balancing test

The interests of the controller or third parties conducting VI/VIN, which involve processing personal data to assess the risk of collisions and prevent harm to individuals and property, may be overridden by the interests or fundamental rights and freedoms of the data subjects. This is because the processing of personal data for this purpose may potentially have negative impacts on the rights and freedoms of the individuals whose personal data is being processed. The following reasons could be argued:

- *The collection and further processing of high-risk personal data, such as logins and passwords*: such data is considered to be particularly sensitive, and its processing may involve significant risks to the privacy and data protection rights of the individuals concerned if not done adequately. In such cases, the interests of the data controller in processing such data for VI/VIN may be overridden by the interests or fundamental rights and freedoms of the data subjects, including their right to privacy and protection of personal data.

- *The sharing or disclosure of personal data with third parties*: data may be shared without the consent or knowledge of the data subjects. This would be the case if a system compromised by attackers led to personal data disclosure or if the operator of the sinkhole server did not understand the data sensitivity and shared logs, for example, with researchers. Indeed, with VIN, to be useful for future analysis, the sinkhole server will need to log the IP address of clients that attempt to connect and the TCP/UDP port numbers of the attempted connections. The scenario is 1) that log information counts as personal data (depending on whether a client IP address will be considered personal data in this context, for example when it relates to a device used by a certain individual) and 2) that an attacker steals that log information. Such sharing or disclosure may be seen as an infringement on the data subjects' right to control their personal data and could potentially have negative impacts on their reputations or opportunities.

- *The processing may be carried out in a manner that is disproportionate or unjustified in relation to the purpose of preventing collisions*: the processing involves collection of potentially sensitive data or extensive amounts of personal data of individuals without their knowledge or consent, whilst the benefits of this method may not justify such large processing in particular if other alternative methods may exist. Indeed, if alternative methods exist that achieve the same or similar purpose with milder means, the processing would not be deemed necessary for pursuing the identified legitimate interest. Therefore, one could not rely on legitimate interests for this reason alone.

Overall, the benefits of conducting VI/VIN may be overridden by the interests or fundamental rights and freedoms of the data subjects if the processing of personal data for this purpose involves significant risks or negative impacts on their privacy and data protection, and potentially other fundamental rights and freedoms and the benefits of VI/VIN are not sufficiently demonstrated to justify such processing.

It's important to note that in all cases, the data controller must ensure that they are processing personal data in accordance with the principles of data protection, including transparency, purpose limitation, data minimization, accuracy, storage limitation, and security, as briefly described below. Additionally, data subjects have various rights under data protection laws, such as the right to access, rectify, erase, and object to the processing of their personal data.

2. Data minimisation

The principle of data minimization is a fundamental principle of personal data protection that requires organizations to only collect and process personal data that is necessary for the intended purpose. This principle is designed to limit the amount of personal data that is collected and processed, reducing the risk of data breaches and privacy violations. In the context of VI/VIN, the principle of data minimization may be challenged because the technique requires collecting and processing data that may not be strictly necessary for the intended purpose. For example, with VIN in particular, the parameters sent in HTTP requests can include sensitive data, such as usernames and passwords.

To address this issue, entities conducting VI/VIN should take steps to minimize the amount of personal data collected and processed. This could include implementing appropriate data protection measures to limit access to personal data, and regularly reviewing and deleting personal data that is no longer necessary or processing the data without retaining it in the first place. For the latter, processing the data would include even receiving the data without retaining it. It is important to note that even if the data is discarded, it is still "processed".

-> Difficulty to implement: medium to high. While it might be possible to conduct VI/VIN without retaining the personal data that might be processed, it remains that an unpredictable amount of personal data would be processed. Implementing the other safeguards to limit the amount of personal data processed would require knowing in advance what data would be collected when conducting VI/VIN, which is impossible.

3. Transparency of the processing

When a TLD is undergoing VI/VIN, the DNS queries from any system that attempts to resolve a domain name under that TLD are ultimately received by the TLD's name servers.[1] This means that information about the domains being queried is processed and this information may include personal data. If the entity conducting the assessment is not transparent about their data collection practices, this could be seen as invasive and raise concerns about data privacy.

The need for transparency of the processing is even more important as there is a risk of disclosure of sensitive/confidential information. Indeed, the VIN sinkhole server will listen on the appropriate TCP port for HTTP requests so that it can respond with a human-readable message explaining the collision scenario and giving more information. However, HTTP requests often include parameters, and some of these parameters could include sensitive information, such as usernames and passwords, for example GET /login?username=janedoe&password=p@ssw0rd!&ip=192.168.1.10 HTTP/1.1.

Entities conducting VI/VIN should be transparent about their data collection practices, how the data will be used, and who will have access to it. They should also clearly communicate their privacy notices and policies to end-users to avoid confusion or concern.

-> Difficulty to implement: high. It seems impossible to provide sufficiently clear, detailed and informed transparency notices to all any/all individuals that may get their data processed in the context of VI/VIN.

4. Security of the processing and risks of data breaches

If the data collected during the assessment is not properly protected, it could be subject to a data breach, which could result in the exposure of sensitive information, such as domain names and IP addresses but also (in some circumstances) a lot more risky personal data such as usernames and passwords, etc. This could lead to identity theft, fraud, or other malicious activities. Processing the data without retaining it (i.e. the data is deleted immediately after it has been written in the system memory) would already mitigate the risk of a breach.

[1] Strictly speaking, not every system's queries are sent to the TLD's name servers: some queries may be answered by a recursive resolver from its cache, though the recursive resolver will have previously queried the TLD's name servers to populate its cache in the first place.