DOI:10.1145/3235764

Vinton G. Cerf

Traceability

At a recent workshop on cybersecurity at Ditchley House sponsored by the Ditchley Foundation in the U.K., a primary topic of consideration was how to preserve

the freedom and openness of the Internet while protecting against the harmful behaviors that have emerged in this global medium. That this is a significant challenge cannot be overstated. The bad behaviors range from social network bullying and misinformation to email spam, distributed denial of service attacks, direct cyberattacks against infrastructure, malware propagation, identity theft, and a host of other ills requiring a wide range of technical and legal considerations. That these harmful behaviors can and do cross international boundaries only makes it more difficult to fashion effective responses.

In other columns, I have argued for better software development tools to reduce the common mistakes that lead to vulnerabilities that are exploited. Here, I want to focus on another aspect of response related to law enforcement and tracking down perpetrators. Of course, not all harms are (or perhaps are not yet) illegal, but discovering those who cause them may still be warranted. The recent adoption and implementation of the General Data Protection Regulation (GDPR) in the European Union creates an interesting tension because it highlights the importance and value of privacy while those who do direct or indirect harm must be tracked down and their identities discovered.

In passing, I mention that cryptography has sometimes been blamed for protecting the identity or actions of criminals but it is also a tool for protecting privacy. Arguments have been made

for “back doors” to cryptographic systems but I am of the opinion that such proposals carry extremely high risk to privacy and safety. It is not my intent to argue this question in this column.

What is of interest to me is a concept to which I was introduced at the Ditchley workshop, specifically, differential traceability. The ability to trace bad actors to bring them to justice seems to me an important goal in a civilized society. The tension with privacy protection leads to the idea that only under appropriate conditions can privacy be violated. By way of example, consider license plates on cars. They are usually arbitrary identifiers and special authority is needed to match them with the car owners (unless, of course, they are vanity plates like mine: “Cerfsup”). This is an example of differential traceability; the police department has the authority to demand ownership information from the Department of Motor Vehicles that issues the license plates. Ordinary citizens do not have this authority.

In the Internet environment there are a variety of identifiers associated with users (including corporate users). Domain names, IP addresses, email addresses, and public cryptography keys are examples among many others. Some of these identifiers are dynamic and thus ambiguous. For example, IP addresses are not always permanent and may change (for example, temporary IP addresses assigned at Wi-Fi hotspots) or may be ambiguous in the case of Network Address Translation. Information about

the time of assignment and the party to whom an IP address was assigned may be needed to identify an individual user. There has been considerable debate and even a recent court case regarding requirements to register users in domain name WHOIS databases in the context of the adoption of GDPR. If we are to accomplish the simultaneous objectives of protecting privacy while apprehending those engaged in harmful or criminal behavior on the Internet, we must find some balance between conflicting but desirable outcomes.

This suggests to me that the notion of traceability under (internationally?) agreed circumstances (that is, differential traceability) might be a fruitful concept to explore. In most societies today, it is accepted that we must be identifiable to appropriate authorities under certain conditions (consider border crossings, traffic violation stops as examples). While there are conditions under which apparent anonymity is desirable and even justifiable (whistle-blowing, for example) absolute anonymity is actually quite difficult to achieve (another point made at the Ditchley workshop) and might not be absolutely desirable given the misbehaviors apparent anonymity invites. I expect this is a controversial conclusion and I look forward to subsequent discussion.

Vinton G. Cerf is vice president and Chief Internet Evangelist at Google. He served as ACM president from 2012–2014.

Copyright held by author/owner.