

Untangling Attribution: Moving to Accountability in Cyberspace

Prepared statement by

Robert K. Knake

*International Affairs Fellow in Residence
The Council on Foreign Relations*

Before the

Subcommittee on Technology and Innovation,
Committee on Science and Technology

*United States House of Representatives
2nd Session, 111th Congress*

Hearing on *Planning for the Future of Cyber Attack*

Room 2318

Rayburn House Office Building
Washington, D.C.

Chairman Wu, Ranking Member Smith, and distinguished members of the House Subcommittee on Technology and Innovation, thank you for the opportunity to discuss the role of attack attribution in preventing cyber attacks and how attribution technologies can affect the anonymity and the privacy of Internet users. In your letter of invitation, you asked me to address the following series of questions:

1. As has been stated by many experts, deterrence is a productive way to prevent physical attacks. How can attack attribution play a role in deterring cyber attacks?
2. What are the proper roles of both the government and private industry in developing and improving attack attribution capabilities? What R&D is needed to address capability gaps in attack attribution and who should be responsible for completing that R&D?
3. What are the distinguishing factors between anonymity and privacy? How should we account for both in the development and use of attribution technologies?

4. Is there a need for standards in the development and implementation of attack attribution technologies? Is there a specific need for privacy standards and if so, what should be the government's role in the development of these standards?

Attributions Role in Deterring Cyber Attacks

Let me begin by stating my view that the utility of deterrence in cyber security may be limited and that the problem of attribution has been over-stated for the high end threats that represent a challenge to our national security. In its classic usage, deterrence is the idea of using fear of reprisal in order to dissuade an adversary from launching an attack. For deterrence to work, it is critically important that we know who has carried out the attack and thus attribution is a central component of deterrence strategy. I believe it may be too broad to view deterrence as a productive way to prevent all kinetic attacks. Deterrence was the central concept in preventing a nuclear exchange between the United States and the Soviet Union during the Cold War. It is not, however, a central part of U.S. strategy to prevent terrorist attacks and its importance in preventing conventional military attacks is more limited than in the nuclear case. During the Cold War, deterrence of the use of nuclear weapons was created through the establishment of "Mutually Assured Destruction" or MAD, in which both the United States and the Soviets understood that any use of nuclear weapons would be responded to in kind. The threat of total annihilation kept both sides at bay. Radar and other warning systems provided the mechanism for attributing any nuclear attack and possession of a second strike capability that could provide a nuclear response even after a successful Soviet launch kept the threat of retaliation credible. Equally important, however, was symmetry.

The Soviets as rational actors did not want to see the loss of their cities, industry, and regime in a retaliatory nuclear strike. As long as we had the ability to hold these assets under threat, a Soviet strike against us would not be to their advantage. Such parity does not exist in cyberspace. Attribution may be a secondary problem to the lack of symmetry. Many countries that possess sophisticated offensive capabilities do not have extensive societal reliance on the Internet or networked systems. If attribution could be achieved, deterrence might not follow because a state conducting an attack in cyberspace, may have little to lose through retaliation. The logical solution to this problem is to threaten retaliation through diplomatic or kinetic means outside of cyberspace, responses that could range from the imposition of sanctions to airstrikes. Thus far, despite the onslaught of attacks in cyberspace, no country has chosen to escalate their response outside of cyberspace. Moreover, it may be difficult to achieve proportionality in response to a cyber attack through other means. Deterrence may simply not be a useful concept to address our current state of cyber insecurity.

If deterrence is to be a central part of our cyber security strategy, I believe it is essential that we can answer three questions: First, what degree of certainty in attribution is necessary to take action? Second, what would that action look like? Third, how will we make potential adversaries understand the answers to these questions prior to an incident so that they will be deterred? To begin, I think it is important to breakdown the attribution problem in cyberspace. There are three broad categories of attack that have their own distinct attribution problem. The first attribution problem, the one on which most attention is focused is the attribution problem for attacks carried over the Internet. These attacks are difficult to deter because of the underlying architecture of the Internet, the lack of security on many hosts, and because the individuals or teams carrying out these attacks can do so remotely, from the safe confines of a non-cooperative country. The second attribution problem is for cyber attacks that are not carried over the Internet. Potentially, many of the most dangerous forms of cyber attacks will be carried out against systems that are not connected to the

internet through other delivery mechanisms including attacks using microwave or other radio transmissions, thumb drives, and other portable media like CDs and DVDs. For these attacks against well-defended military and industrial systems, the attribution problem is similar to the attribution problem for kinetic attacks and can be addressed through real world forensics, investigation, and intelligence. Finally, there is the problem of attribution for the introduction of malicious code in the supply chain for hardware and software. The threat to the supply chain may be the area of most concern today, yet the attribution problem for the insertion of malicious content into software and hardware is no different from a traditional investigative challenge to identify the opportunity and the motive for inserting malicious content (see Figure 1 for a visual representation of these challenges).

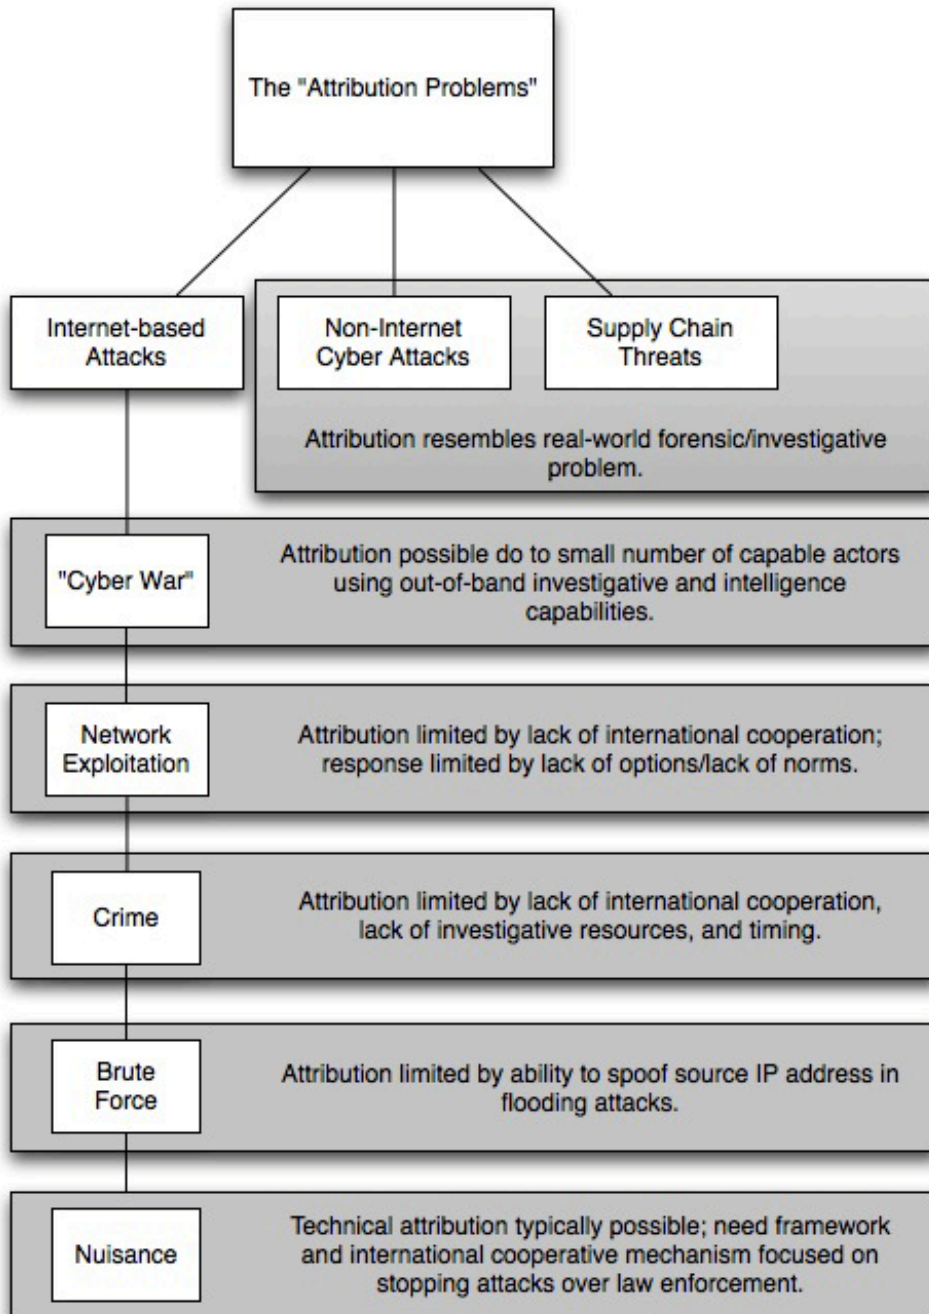


Figure 1: The Attribution Problems

With the exception of flooding attacks, all other forms of Internet-based cyber attack require two way communication between the attacking computer and the victim computer. Sophisticated adversaries will take steps to obfuscate their true location and identity through the use of proxy systems, whether they are compromised computers or anonymization services or both. Despite these precautions, trace back techniques and digital forensics can provide the technical means to allow the attackers to be discovered. The barriers to the use of these techniques are more legal than technical, due to international boundaries and non-cooperative countries. If we breakdown the various threats carried over the Internet, the scope of the attribution problem can be brought into focus and different solutions for managing each threat begin to emerge.

Attacks can be divided into the following categories ordered by the threat they pose: cyber warfare, cyber espionage, brute force attacks, crime, and nuisance. For each of these, both the attribution problem and the issue of response are different. For the highest level threat, that of cyber warfare, the attribution problem is largely overstated. As with other Internet based attacks, technical attribution may be difficult and the forensics work will take time, but at present there are a limited number of actors that are capable of carrying out such attacks. Moreover, the resources, planning, and timeline for such attacks would provide many opportunities to identify and disrupt such attacks. Estimates vary, but on the low end, many experts believe that only four countries possess the capability to carry out a catastrophic attack in cyberspace, the so-called Cyber Pearl Harbor, Cyber 9/11, or Cyber Katrina. On the high end, up to 100 state actors and private groups closely affiliated with state actors may have the capability. No matter which estimate is accurate, this is a fairly small list of suspects that can be narrowed down through technical means, as well as out of band methods that include intelligence, analysis of capabilities and analysis of intent. If not already a priority, US intelligence agencies should be focused on identifying actors with high-level capabilities and understanding their intentions. While it has become a truism that hacking tools can be downloaded off the Internet and used by an individual with little or no technical skills, these tools do not pose the kind of threat that could cause widespread destruction. If the operators of critical systems cannot defend against such attacks, they are not taking the threat seriously. As the relevant technologies continue to evolve, it is important that the difficulty in carrying out significant attacks increases. Our critical industries, military and government agencies must continue to raise their defense levels in order to keep the ability to cause destruction in the hands of a limited number of state actors.

In the event of a catastrophic cyber attack, attribution to at least some level will almost always be possible. The question becomes to what level of certainty must attribution be demonstrated in order for the President to take action? At the lowest level, attribution that traces an attack back one hop can provide the foundation for further investigations. If that first hop is in a non-cooperative country that is unwilling to assist in the investigation, that may be enough evidence to hold that country accountable. As with the 9/11 attacks when the Taliban refused to turn over Osama Bin Laden, it may be appropriate under such circumstances to hold a non-cooperative country accountable, a concept I will return to later in this testimony.

On the issue of espionage, the capability necessary for network exploitation is generally lower than that required for destructive attacks, particularly in the realm of economic espionage where private sector companies are targeted. What we lack is not so much an ability to attribute attacks, but international norms that keep espionage limited. Espionage is generally recognized to be permissible under certain circumstances and many scholars will argue that it has a stabilizing effect on the international system by reducing paranoia. As has been recently demonstrated by the discovery of a Russian spy ring in the United States, engaging in

espionage is not necessarily considered a hostile act and can be resolved without further escalation. The challenge with cyber espionage is that we lack norms that limit the extent to which states engage in it. This problem is exacerbated by the fact that cyber espionage is not constrained by the costs, consequences and limitations of traditional espionage.

By way of example, consider the case of Robert Hanssen, a former FBI agent who spied for the Soviets and then the Russian Federation for over two decades. Over that period, Hanssen smuggled several hundred pages of classified material to the Russians, who paid him several hundred thousand dollars and maintained a network of handlers in order to run this operation. Hanssen paid a heavy price for his betrayal. Having been sentenced to life in prison, he spends 23 hours a day in solitary confinement at a Supermax Facility and is addressed by the guards only in the third person (“the prisoner will exit the cell.”) The American spies he betrayed inside Russia were not so lucky. Most were executed. During the Cold War, spying had consequences. Now, according to public media reports, foreign intelligence agencies have exfiltrated several terabytes of information from US government systems.

Whatever country or countries are behind this espionage campaign, the people who are carrying it out are working safely from within the borders of their own country at little risk of being discovered or imprisoned. The low cost and low risk of cyber espionage is the problem, not the difficulty in attributing the source of the activity. If ironclad proof emerged of who was behind an incident of cyber espionage, what would the U.S. response be, particularly given the likely intelligence advantages that the United States gains from cyber espionage? It may be time that we recognize cyber espionage to be a different phenomenon from traditional espionage, one that requires a different set of norms and responses. I doubt, however, that we lack sufficient certainty of who is behind these campaigns that we are limited in our response simply because we do not know who is carrying them out.

Brute force attacks, so called distributed denial of service attacks or DDOS attacks, do present a specific technical attribution challenge. During these attacks, compromised systems formed into a botnet flood targets with large numbers of packets that do not require the targeted system to respond. The malware behind these attacks will provide false information on the source of the packets, so that the machines sending the packets cannot be identified. This particular problem is due to the trusting nature of the internet protocol which does not provide any security mechanism to keep this information from being falsified. To deter DDOS attacks, it may be necessary to strengthen the Internet Protocol so that attacks can be traced to the computers that are part of the attacking botnet, and from their to the command and control servers and potentially to the botnet master himself. It may be equally productive to simply locate compromised computers participating in the attack and shut these down.

For crime, the goal of attribution is to aid in investigation and result in criminal prosecution. Attribution is therefore necessary in the first instance to direct where an investigation should be targeted and for this first step, attribution needs to rise to the level sufficient for ‘probable cause’ to initiate the investigation. This first level of attribution may only need to lead to a system, not to an individual and an IP address is often times all that is sufficient. In turn, the investigation will need to establish attribution to an individual or group of individuals for the purpose of prosecution. For prosecution to be successful, attribution will need to rise to the level of guilt beyond a reasonable doubt. In between, there is the potential to pursue criminals through civil litigation, in which case the standard for attribution would be lower, and guilt would be assigned based upon a preponderance of the evidence. The problem is that currently, many countries lack both the legal framework and resources to pursue cybercrimes committed by their citizens or that use systems within their territory

that target victims in another country. Even crimes committed by individuals in the United States against individuals in the United States will make use of intermediary systems in other countries, particularly those that are not likely or able to cooperate with an investigation. What is needed to deal with the problem of crime is not better attribution but stronger legal mechanisms for working across international borders, the ability to shutdown attacks as they are taking place, and more investigative resources. Ultimately, there must be penalties for states that do not cooperate in investigations and do not take steps to secure their portion of cyberspace.

For nuisance attacks, attribution is rarely a problem. The problem is that few if any investigative resources are assigned to cyber criminal activity that does not have a high monetary value associated with it. This is a situation in which the impact of the crimes committed is fairly low but the resources necessary to address them are high given the volume of the problem. As an example, look at the problem of SPAM. The 2003 CAN-SPAM Act requires spammers to provide accurate header information and to provide an opt-out method for recipients so they can choose not to receive future methods. Yet nearly a decade later, SPAM is flourishing as 9 out of 10 emails are SPAM. For most of these messages, the organization that sent the message is identifiable because they are selling a product. What we lack is an enforcement method that fits this problem, one that is focused on stopping the nuisance behavior rather than prosecuting those who are behind it. Similarly, nuisance level network attacks, the type that can be initiated through downloads off the Internet, are rarely investigated and prosecuted yet they distract system administrators and computer response teams from higher level threats. Investigating and prosecuting more of this behavior could deter many of the people who engage in it.

For most of these threats, the challenges are not so much related to attribution as they are to resources and international cooperation. Focusing on deterrence may simply be the wrong way to think about how to handle these problems. The threats are materializing every day, making the abstract theorizing that laid the foundation for deterrence in a nuclear confrontation unnecessary. They are also, in every respect, a lower level concern that in no way threatens the existence of the United States. Instead we should focus in two areas. We need to reduce the scale of the problem by stopping threats as they unfold and by reducing the vulnerabilities that the threat actors make use of in their attacks. An investigative and enforcement approach to all problems is simply not tenable. Instead of trying to trace every incident back to a human user, we need to develop a legal framework for stopping attacking systems. We must move beyond treating intermediary systems as victims, and start viewing them as accomplices. In the United States, such a framework could require ISPs to monitor their network for compromised systems that have become parts of botnets and quarantine those systems until the problem is resolved. Similarly, we need mechanisms that allow companies or individuals that are under attack and have traced the attack to a system or systems to request for those systems to be shutdown. This process needs to take place quickly and mechanisms must be developed to authenticate such requests across international borders. Such a framework, if developed in the United States, could be promoted as a global model.

For higher end threats, there are lessons we can learn from the last decade of dealing with terrorist threats. The key is to move beyond the search for perfect attribution and instead hold states that do not cooperate accountable. Currently, the situation can be summed up like this. When an attack is traced to another country that is not cooperative, the investigation dead ends. If that country is Russia, Russian authorities will typically say that the incident was carried out either by patriotic hackers or cyber criminal groups that the Russian government cannot control. If that country is China, Chinese officials will point out that China is often the

victim of cybercrime and that do to the poor security on many Chinese systems, they are often compromised in an effort to cast blame on China. In both cases, national sovereignty will be raised to explain why cooperation cannot be more forthcoming.

To move beyond this stalemate, the United States should make public a position that treats failure to cooperate in investigating a cyber attack as culpability for the attack. Countries should know that they can choose to have the incident treated as a law enforcement matter by cooperating in the investigation or choose not to cooperate and have the incident treated as a hostile attack for which their country will be held accountable. Over the last decade the concept of state sovereignty has evolved so that sovereignty not only comes with rights in the international system but also responsibilities. The evolution of this concept is due to events in one of the least wired parts of the world: the Hindu Kush.

In 1999, Michael Sheehan, the U.S Ambassador at Large for Counterterrorism delivered a demarche over the phone to the Taliban's foreign secretary. The message was clear: as long as the Taliban continued to harbor and support al Qaeda and its leaders, the United States would hold the Taliban responsible for any al Qaeda attacks against the United States or other countries. To drive home the point, Sheehan used an analogy. He told the Taliban's representative: "If you have an arsonist in your basement; and every night he goes out and burns down a neighbor's house, and you know this is going on, then you can't claim you aren't responsible." The United States made good on Ambassador Sheehan's word after 9/11, and as the international community attempts to address failed states that cannot control their borders or police their internal territory, this new concept of sovereign responsibility is taking hold.

Applying this new concept of sovereignty to cyberspace has its merits. As with al Qaeda in Afghanistan, failure of a state to prevent its territory from being used to stage an international cyber attack should not, in and of itself, constitute a violation of state responsibility. Indeed, a world in which states monitor and constrain citizen activities to prevent crimes before they take place would be a very frightening world. What is crucial, however, is how states respond when confronted with the use of systems within their territory for cyber attack. If the Taliban had responded to requests to turn over bin Laden, the invasion of Afghanistan might never have occurred. Based on this new paradigm of sovereignty, states should be expected to pass laws making international cybercrime illegal and enforce them. They should have mechanisms in place to respond to international requests for assistance and they should have some ability to oversee the hygiene of their national networks. Better attribution through post-incident forensic techniques will be a crucial part of this new paradigm, but the development of ironclad attribution, will not necessarily lead to better security in cyberspace.

The Role of Government and Private Industry in Improving Attack Attribution

In order to improve attack attribution, there are many things that can be done with current technology. The most crucial is for both government and private industry to do a better job detecting significant threats, mitigating them quickly, and capturing evidence that can be used by law enforcement for investigative purposes. Forensic techniques are getting better, but there are genuine civil liberties concerns with them getting too good.

The vision of perfect attribution can best be summed up as the idea of giving packets license plates. Under such a system, compromised systems or other proxies could not be used to hide the identity of attackers because each packet would be labeled with a unique identifier, possibly an IPv6 address that has been

assigned to an individual after having that individual's identity authenticated in some verifiable way. Access to the network would require authentication, and each packet produced by the user would be traceable back to that user. The privacy implications of such a system would be obvious, turning the Internet into the ultimate tool of state surveillance. The security benefits for pursuing criminals and state actors, however, would be minimal. Without cooperation from all foreign states, criminal activity will simply gravitate to states that do not authenticate identity before issuing identification numbers or choose not to participate in the system at all. Many states benefit tremendously from cybercrime, both directly through the cash it brings into economies, and indirectly through the bolstering of technology development through the theft of intellectual capital. Moreover, for less capable states, cybercrime provides the necessary cover of darkness for espionage to take place. By cracking down on cybercriminal groups, the activities of state actors would stand out starkly. Ultimately, such a system would restrict the freedom and privacy of most users, while doing little to curb criminal elements or state actors who would find ways around the system.

As a baseline, of what we should expect from digital forensics, it may be instructive to look at the role forensics plays in the real world. Many people have become familiar with modern forensics techniques through the popular series CSI and its spinoffs, television shows about real-world crime scene investigators. Each episode begins with a body. The crime scene investigators come in and walk the scene collecting forensic evidence and then take it back to the lab and process it for clues. This activity takes us to the first commercial break in an hour-long drama. The forensics have yielded clues about who the victim was, how he or she was killed, and possible attributes of the killer. Then the detective work begins. The detectives try and establish a motive. They delve into the past of the victim. They ask themselves who would have wanted the victim dead? They ask a lot of questions of a lot of people. On television, this process is packed into an hour. In the real world it can take days to weeks, months and years.

Cyberspace isn't so different from the real world. We have digital forensic tools and trace-back techniques that in the latest incident with Google, allowed the company to conclude that the attacks emanated from China. We can't know more than that without some good old-fashioned investigative work but we can ascertain motive based on what systems were infiltrated and what data was stolen. We can narrow down the list of possible suspects by geography. We can further narrow down the set by capability. Only so many people in the world have the ability to put together the kind of code used in the hack. We also know whoever built the exploits wasn't working alone. That's enough leads to get an investigation going in the real world, and it is also enough in cyberspace.

While the Google case illustrates the attribution "problem", it also illustrates the need for Internet Freedom, something the Chinese government is trying to erode. Our law enforcement community might want ironclad attribution on the Internet to combat cyber crime, but the Chinese government and other authoritarian states want it to combat speech. We may want to know who carried out the hacking of Google but we also want to protect the identity of anonymous posters in online forums about Chinese human rights.

Creating the perfect surveillance state online is within our technical means. In real-world equivalents, we could label each packet with its digital DNA, tying it to a single real-world person, and recordings of everything that goes on so we can play back the tape. But cyberspace isn't so different from the real world, especially since more and more of what we used to do by walking we now do online. If we don't want to live in

a surveillance society out here, we also do not want to live in one in cyberspace. The tools for digital forensics are getting better. We don't want them to get too good. What the Google incident really demonstrates, isn't a technical problem; it's a legal and diplomatic one. We lack norms for acceptable behavior by states in conducting espionage online and we lack agreements between states to partner in pursuing cross-border cyber criminal activity. Better surveillance wouldn't solve that problem.

In two narrow areas, government and private sector technology companies should collaborate to improve two of the basic protocols that govern internet transactions. First, government and industry must work together to develop a secure version of the basic internet protocol that authenticates the "from" information contained in packet headers. In distributed denial of service or DDOS attacks that do not require the return of information, the ability to supply false sender information makes it difficult to trace and block such attacks. Similarly, the underlying protocols for sending email allow an individual to spoof the identity of a sender so that someone with malicious intent can send email appearing to be from a bank, a friend, or a work colleague. This weakness is typically exploited in social engineering attacks in order to get the recipient to click on a link that will download malware or send back sensitive information. These problems are well known and well documented. After more than two decades, I believe it is safe to conclude that the informal, consensus-based processes used by the Internet Engineering Task Force to develop and adopt new protocols will not solve these problems. The federal government must step in, lay out the challenge, and lead the development and adoption of protocols that solve these problems. An "X-prize" strategy might prove useful in this context.

Privacy and Anonymity in Resolving Attack Attribution

In the early days of the Internet, anonymity was how privacy was obtained when online. As a general trend, anonymity on the web is eroding for most users due to the interactive nature of current web content but new ways of protecting privacy have not developed, at least not for the average user. In terms of protecting privacy, anonymity is only useful in a "web 1.0" context. In the web 1.0 era, users were passive recipients of information posted to the web. Anonymity on the web is still useful for accessing information that you do not want others to know you have accessed, whether it be pornographic material or information on democracy if you live under an authoritarian regime. Increasingly, however, access to information is not what the Internet is being used for. Managing health records and finances and communicating online cannot be done anonymously. What is needed is privacy, something that does not currently exist on the web that must be created through both technical and legal mechanisms.

Most of the so-called "free" web is funded through advertising, and advertising is increasingly targeted to individuals based on information collected about them from their IP address and from various types of cookies placed on their computers when they access sites. By the time my homepage at the nytimes.com has loaded, a total of 12 cookies have been loaded onto my computer, including "flash cookies" that cannot be deleted through standard browser settings. While some of these cookies are used to authenticate my username and password on the site, the vast majority are for advertising, meant to track my use of the internet in order to target advertising at me. Companies sell geo-location services that use IP information to determine where you live so that advertising can be targeted at you for local services. By default, my browser, my computer, and the

websites I visit are set to allow all this to happen without me knowing it. Advanced users may have the skill set and the motivation to set their browser settings and take other steps to avoid privacy loss but most users do not.

At present, only the technically sophisticated, be they law-abiding citizens concerned with their civil liberties or criminal actors, can obtain anonymity, while the average Internet user experiences a total loss of privacy. As the technology develops to improve attribution, we need to ensure that our laws develop to protect their use, both by government and by the private sector. These points to the need for government intervention to require companies that collect information online and track users to be explicit about what they are doing. Surrendering your privacy online in exchange for “free” access to information should not be something that happens behind the scenes, but an explicit decision that users make. The equivalent of the Surgeon General’s warning, something short, explicit, prominent and standard should be displayed on sites that use privacy compromising methods to generate advertising revenue.

In order to protect private communication online, we need to implement both technical solutions and stronger legal protections for the content of communication. While law enforcement and intelligence agencies are restricted from accessing private information without due process, private sector entities and criminals have far fewer barriers. The average home users email messages are not secured end-to-end through encryption, and the laws that protect the intercept of these messages are far weaker than those that protect regular mail.

Taken together, these steps would replace the loss of anonymity that was the foundation of privacy on the early web, with privacy for all activities carried out over the Internet, including transactions and two-way communication.

Standards Development for Attack Attribution and Privacy

As stated previously, I believe it is necessary for the US government to work with the Internet engineering community to address known problems in the current suite of protocols. In my view, these problems are both limited and correctable but both funding for development and incentives for adoption post-development are necessary. The goal should not be to create ironclad attribution that would turn the Internet into the ultimate tool of the surveillance state. Rather, the end state should be protocols that prevent the spoofing of IP addresses and email.

On privacy standards, I believe that it is government’s role to protect the privacy of individual users. Government must stop assuming that consumers have all the information they need to make informed decisions about privacy. The goal of government intervention in this area should be to make the decision to surrender privacy in exchange for access to information and services a transparent decision. Websites should be required to notify users if access requires the installation of cookies that will track users for the purpose of targeting advertising. Many if not most users may make the decision to surrender their privacy for access to so-called “free content”. Others may choose a pay option. Still others may seek out content that neither costs privacy or dollars.

These two issues overlap for Internet Service Providers. The activity of ISPs is largely unregulated in the United States. For ISPs, attribution on their networks is not a problem: they can see malicious activity and trace it back to a customer. When evidence of the next jump on a host has been deleted, ISPs are often able to trace the next hop of packets. Standards are necessary for what ISPs should and should not be required to

track, for how long they should store such information, and how this information can be shared with law enforcement or private parties.

Finally, we need standards for the operation of anonymity services. Services like Hotspot Shield, Tor, and others provide a valuable service to many Internet users, particularly those living under authoritarian regimes where accessing certain websites may not be possible or may be tracked in order to identify dissidents. Yet these same systems can be used for criminal purposes. Standards are necessary for regulating these services and they must be promoted internationally. These services provide anonymity, which, as previously discussed, is only useful for accessing information sources and anonymous posting activity. These services should therefore restrict their users to web-based activity. They should also make it easy for companies and government agencies to block the outbound IP addresses to prevent users that have gained anonymity from attempting to access secure systems. If you are trying to access your own bank account online, there is no legitimate reason to use an anonymization service. Finally, these services should retain auditable logs for law enforcement purposes. Users should understand that this information will be kept private, and only released if the service has been used for criminal purposes. Ultimately, as with states, anonymization services should be held accountable for their users' behavior if they do not cooperate with law enforcement.

Conclusion

As I have expressed throughout this testimony, it is my view that the problem of attribution has been largely overstated. Ironclad or perfect attribution would not address the problems of cyber warfare, espionage, crime or other threats in cyberspace. Such a capability would, however, be injurious to freedom of expression and access to information for many people around the world. Stronger mechanisms for international law enforcement cooperation are necessary, as is the ability to stop attacks in progress, and improvements to the general hygiene of the Internet ecosystem. More than anything else, we need to develop better and stronger options for responding to threats in cyberspace and introduce consequences for states that do not cooperate in stopping attacks or in investigating them. Finally, we need to move beyond anonymity as the guarantor of privacy on the Internet and instead work to create privacy through both technical means and legal requirements. Thank you for the opportunity to testify on these important issues. I would be happy to answer any questions at this time.

Robert K. Knake

Robert K. Knake is an international affairs fellow in residence at the Council on Foreign Relations studying cyber war. He is currently working on a Council Special Report on internet governance and security. Prior to his fellowship, he was a principal at Good Harbor Consulting, a security strategy consulting firm with offices in Washington, DC; Boston, MA; and Abu Dhabi, UAE, where he served domestic and foreign clients on cyber security and homeland security projects. Rob joined Good Harbor after earning his MA from Harvard University's Kennedy School of Government. He has written extensively on cyber security, counterterrorism and homeland security issues. He is co-author (with Richard Clarke) of *Cyber War: The Next Threat to National Security and What To Do About It* (HarperCollins, April 2010).