The ISP Column
A monthly column on things Internet

October 2013
Geoff Huston

The Big Bad Internet

I often think there are two types of stories about the Internet. One is a continuing story of prodigious technology that continues to shrink in physical size and at the same time continue to dazzle and amaze us. We've managed to get the cost and form factor of computers down to that of an ordinary wrist watch, or even into a pair of glasses, and embed rich functionality into almost everything. The other is a darker evolving story of the associated vulnerabilities of this technology, where we've seen "hacking" turn into organised crime and from there into a scale of sophistication that is sometimes termed "cyber warfare". And in this same darker theme one could add the current set of stories about various forms of state sponsored surveillance and espionage on the net. In this article I'd like to wander into this darker side of the Internet and briefly look at some of the current issues in this area of cybercrime, based on some conferences and workshops I've attended recently.

There is little doubt that the Internet has been populated in many diverse ways, and at the same time as we see a proliferation of online services for users we also see proliferation of attacks on these same services and users. Some attackers want to relieve a victim of their money, while others attackers head into areas of disruption, intelligence collection and espionage. No doubt there are many other motives at play as well. At the same time as the variety and efficacy of such attacks escalate, there is a level of frustration that our legal framework is not keeping up to date in classifying such acts as criminal activities, and that our law enforcement and intelligence agencies are unable to undertake their roles and offer us those basic protections under law that we have come to simply assume in the physical world.

Just how different this cyber world can be was highlighted in a recent presentation at a cybercrime conference from Fox-IT, the private Dutch organisation that, notably, performed the forensic examination of the consequences of the Diginotar certificate compromise. This presentation started with the proposition that, in the area of cyber crime, LEAs are effectively blind. They are unresponsive to small scale issues, and while this may lead to a call for increased presence and surveillance in some quarters, there is also the sensitivities of state-based online surveillance that hinder acceptance of the value of such increased LEA surveillance online. And while larger organisations have the wherewithal to operate a competent security function, public users enjoy no such facility. The legislative framework is dated, given that it is constructed on a premise of a physical world, physical criminal actions and physical evidence. In addition, there is an acute skills crisis in this sector. These considerations have lead to the rise of private cyber investigation entities, such as Fox-IT, Crowdstrike and Mandiant. The motivations of the private entities in this space are quite different to the LEAs. The victim is the focal point, not the crime, the criminal act is more important and the underlying trend, and catching the perpetrator(s) is less important than mitigating the damage caused to the victim and preventing its recurrence. The private investigation is not hampered by physical borders and does not necessarily require physical presence. Are private investigations more effective than traditional public policing? Probably not. Indeed its probably the case that their motivations are different enough that direct comparisons are not all that useful. Online criminals run a minimal risk of discovery and capture from private investigators. as these private investigators generally have no overarching motivation to catch the criminal per se. This leads to a problematic mismatch in this space. Our common dependency on the Internet continues to grow, and grow rapidly, but our security capabilities lag far behind, and this lag is getting worse. In the public sector legislation, skills and methods all need attention. The private sector response has filled the needs of some individual entities who can afford to use their services, but such activities do not address the broader aspects of the vulnerabilities in our environment that facilitated the criminal act, nor are they focussed on the apprehension of the criminal.

A related story comes from McAfee, which takes the position that these days cybercrime is a service-oriented business enterprise. The presentation commenced from the notion of the criminal enterprise as a service enterprise where competitive pressures are dropping the barriers to entry, such that the instruments of an effective cyberattack can be outsourced completely. The presentation pointed out the capability to purchase online zero day vulnerabilities, email lists, and target addresses. In terms of cybercrime product development you can now purchase exploits, you can purchase the testing of malware against current anti-virus software, the service providers in this activity now gather reputation and credibility in the same was as online vendors in eBay. Some spam providers evidently provide a chat window and a 24 hour help desk. In this highly competitive market for cybercrime services there is also a price slump, and bot armies can be leased for nominal sums. Similar services exist in the hacking space, where email passwords can be cracked for a fee, as can credit cards. This is spreading into currency, where Bitcoin is evidently widely associated with cybercrime. The volume of this activity is increasing, the severity is increasing, and the LEA response is largely non existent. (http://www.mcafee.com/de/resources/white-papers/wp-cybercrime-exposed.pdf)

Trend Micro have looked at the pervasive use of inline apps and services, and the scale of the risk that cybercrime represents to our society, and produced a work of what one could call either speculative near future reality TV, or a webcast drama on the theme of the pervasive nature of the generation and use of personal information and the potential ways this could be perverted. (http://2020.trendmicro.com) Microsoft operate a data collection subsystem within a larger framework called "Operation b54", which is being used as a means to gain some real time visibility into botnet formation and operation. At the same time as the botnet proliferation, Microsoft report that they are seeing the increased use of TOR, Bitcoins, and VPS solutions where the technically savvy users are adept at burying themselves and making online anonymity commonplace.

The Law Enforcement Agency (LEA) sector is facing considerable challenges at this point in time. One view from within the LEA sector is that the level of this criminal activity is escalating at unprecedented levels, while the resources and skills available to the LEAs continue to be woefully inadequate. The LEAs appear to be lagging behind the privateers, and the privateers are rapidly innovating at a pace commensurate with their criminal counterparts, which keeps LEAs constantly back-footed, despite protestations to the contrary. There is a similar story with some of the Computer Emergency Response Teams (CERTS) and the LEAs, where there are claims that the relationship between the CERT teams and the LEAs is not working as well as it could, and we hear from some CERTS that the LEAs are unresponsive and not overly cooperative, and a similar claim is heard in the other direction. Again, there is a difference of motivation, where the CERT is strongly motivated to assist its clients, the potential or actual victims, while the LEA focus is often directed to the crime itself and its perpetrators.

In the meantime, over on the Internet, things are not exactly getting any better.

These days the current picture of SPAM is that it overwhelms the "real" email traffic. Some spam fighting groups have claimed that spam outnumbers genuine mail by a factor of 100:1. Whatever we could claim as the success of the various responses of blacklists, mail filters, reputation services, certification and regulation, the basic observation is that the escalating volume of spam has managed to readily out scale the effectiveness of the responses. At the same time we've now managed to imbue IP addresses with "reputation", and its difficult to understand whether this a a net positive, as the action certainly has its downsides. Many readers would be aware of the various forms of blacklists that list the "disreputable" IP addresses. These blacklists enumerate the IP addresses of hosts that have been observed to emit spam, control a botnet, host phishing web sites, or any of a number of related nefarious activities. Once an IP address is listed in one of these lists, then many other systems will not communicate with it. It’s often claimed that it’s extremely easy to get an IP address into one of these blacklists, but very hard to get it off them once it has been listed. Jane Austen may well have been talking about IP blacklists when in Pride and Prejudice she had Mr Darcy confess that "My good opinion once lost is lost forever." Part of the problem here is that it's very easy to set up a blacklist, as many folk have already done. So if an address has gained a bad reputation, understanding the context and understanding where it has gained this blacklist status is often a challenge. You can tell when a concept has gone perhaps a little too far when aggregators enter the fray. As is claimed on the dnsbl.info web site: "DNSBL Information provides a single place where you can check that status of your mail server's IP address on more than 100 DNS based blacklists."

The fact that software has its vulnerabilities should come as a surprise to no one. And the observation that such vulnerabilities have been exploited for criminal acts should also not really come as a surprise. So we should fix these vulnerabilities – right? However, sometimes this decision to fix the vulnerabilities in a deployed system transforms itself from a single operation of finding all vulnerabilities, patching the code and circulating the updates, to that of repeated iterations of collecting all currently known vulnerabilities, patching and circulating, and for some systems this process persists for years and becomes a task that is seemingly without end. The longer the software is deployed and the larger the population of deployed systems, the higher the incidence of detected vulnerabilities. Windows XP falls into that category of a widely deployed and now quite a venerable platform. XP appears to account for around a third of the currently deployed desktop and laptop systems in today's Internet. (http://en.wikipedia.org/wiki/Usage_share_of_operating_systems) Microsoft has released a large number of patches for this system over the years, yet the vulnerabilities in the deployed population appear to persist. Worryingly, Microsoft have announced that further support for Windows XP will cease as of 8th April, 2014. (http://windows.microsoft.com/en-us/windows/end-support-help) To further complicate the issue, it appears that unlicensed copies this particular operating system have been widely distributed over the years, and with the centralised form of update management used by Microsoft, it appears that many users of these pirate copies of the system believe that applying updates to their system will invalidate their system. So we see continuing use of original vulnerable systems on an Internet where the volume and sophistication of the attack probes overwhelm the residual defences of these old unpatched systems. The persistence of so-called botnet armies of corrupted systems that are under remote control is one of the more disturbing outcomes of this situation, but the larger picture relates to the increasing dependency that we place on the networked environment as part of our lives, and the concerns that many of the elements that support this environment use neglected and highly vulnerable software. From the signs at airports, to cash terminals at retail outlets, to thousands of other deployments that range from the mundane to the vital, it appears that XP is still prolific, and its vulnerabilities are a source of serious concern.

But it's not all a case of software vulnerabilities is end systems. We also see instances where the IP protocols are turned against us. One of the more effective denial of service attacks on today's network is a DNS reflection attack, where DNS resolvers can be turned into uncontrolled traffic generators, with a result that is capable to swamping many service providers with this imposition of unsolicited and unwanted traffic. We would not be so vulnerable to this form of attack if it were so easy to pass packets through the networks with a crafted source address. But our efforts to convince the network operators that everyone benefits if they maintain so-called BCP-38 filters on their outbound traffic has largely fallen on deaf ears. The document, BCP-38, was published 13 years ago, as RFC 2827 in May 2000, so no one could really claim that this is too recent so they haven't had the time to get around to implementing it yet. This is an instance of a more general observation about our behaviour: when the assessment of the risk of an event occurring, even when multiplied by the assessment of the directly incurred damage that may result from such an event, is less than the marginal cost of mitigation, we simply don't act to try and reduce the level of exposure to risk. As long as these network operators believe it to be highly improbable that they themselves will be the targets of such an attack, or if the cost of such an attack is relatively minor, then they see little reason to spend money to mitigate the risk to their customers by maintaining rudimentary source address filters that radically narrow the scope of such attacks that rely on the ability to perform source address spoofing.

Can't we use all these clever networking technologies to track down all these network nasties and turn them off? There is a widely held impression, reinforced by the PRISM stories, that the online environment is one that admits little in the way of personal privacy and true anonymity. One is led to believe that the thick plumes our of digital exhaust are carefully stored and analysed. No deed goes unnoticed, and no act is truly anonymous. But this is probably a somewhat mistaken impression. Depending on the networks involved, the technologies used in the networks, whether or not the networks even perform rudimentary logging, and the attendant the issues of correlating various logs of all these activities, whether nefarious or not, then the true ability to perform such forms of extensive tracing and tracking is perhaps context dependant. The result is a network that appears to be an eclectic mixture of a set of fish bowls and dark alleyways, without any real way of being able to figure out in advance where precisely we are at any point.

Law enforcement agencies are exposed to the same variability, which implies that in many ways their ability to respond to cyber crime in ways that match society's expectations varies. Sometimes bad acts on the Internet are readily exposed, and the criminal perpetrators along with it. But other times we find LEAs they are lacking the essential skills, resources and basic forensic data to respond in a meaningful way, and the private investigators are there to fill the gap.

The network is truly a place that has its dark and hostile corners and many bad deeds not only go unpunished, but often go undetected by all but the victim and the perpetrator. And, on the whole, for a vital public communications utility in today's world, that's probably not a very reassuring place to find ourselves.

Disclaimer

The views expressed are the authors' and not those of APNIC, unless APNIC is specifically identified as the author of the communication. APNIC will not be legally responsible in contract, tort or otherwise for any statement made in this publication.

About the Author

Geoff Huston B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. He is author of a number of Internet-related books, and was a member of the Internet Architecture Board from 1999 until 2005, and served on the Board of Trustees of the Internet Society from 1992 until 2001.
www.potaroo.net