serene risc Digest Volume 2, Issue 1

Cutting Edge Research Summaries for Policy-Makers and Practitioners

There are many phishes in the sea, but are they all unique?

Phishing is a computer security threat that involves messages designed to get people to do things for the benefit of an attacker. It is important to properly understand phishing messages to effectively identify and avoid them. Cui and collegues from industry and academia collected nearly 20,000 phishing websites over 10 months. Instead of comparing phishing websites to legitimate websites, they compared the content structure of phishing websites. The researchers found that 90% of phishing websites were repeats of an earlier attack. In order to avoid detection, attackers often make subtle changes, such as switching the page to a different domain or subdomain. They seem less likely to change the structure or content of the pages. This means that although many phishing websites do not last for long, certain content structures may be recycled and reused over longer periods. Detecting phishing based on content structure could identify a large proportion of phishing websites. Requiring attackers to modify the structure of their phishing pages could slow down their rate of adaptation to upcoming filtering technologies.

Does click fraud have a weak link?

Online advertising helps facilitate the development of the Internet, as it provides revenue for content creators. Organizations often pay advertisers for the numbers of views or clicks on an advertisement assuming it will lead to sales. However, fraudsters can generate fake traffic imitating human activity, without leading to any real revenue. This is known as "click fraud." They often use "click-fraud botnets" to carry out their fraudulent schemes. A click fraud botnet is a network of hacked computers that generates fakes traffic. Faou et al. looked at a sample of traffic generated by a well-known click fraud botnet to build a picture of the malware-generated advertising traffic. This allowed them to find and map the relationships of 225 actors potentially involved in the network. By looking at the ties between the actors, the researchers were able to identify the key players of the network and find whose removal from the network would disrupt it most. However, real word practicalities can make removing some actors difficult. It is more feasible to remove or isolate the more legitimate companies from the botnet network. Encouraging these businesses to disconnect themselves from fraudulent operators may also be a longer-term click fraud botnet takedown strategy.

Do industry qualifications test competency effectively?

Cybersecurity competency is often expressed in terms of qualifications. However, the methods used to assess competency vary between qualifications. Knowles et al. reviewed 74 industry cybersecurity qualifications to see which methods they use. They found five distinct methods: multiple choice exams, narrative form exams, oral exams, virtual labs and employment reviews. Low-cost multiple-choice exams were the most popular method, present in over 80% of qualifications, and were also the sole form of assessment for almost half of all qualifications. The researchers also surveyed 153 various industry stakeholders to see how effective they perceived each testing method to be at measuring cybersecurity competence. Participants did not seem to find multiple choice exams to be very effective, with nearly half perceiving it as a fair or poor competency assessment method. Although costlier, they favoured employment reviews and virtual labs. It seems that a large proportion of cybersecurity industry qualifications use methods perceived as being ineffective. Understanding the appropriateness of testing and the relative costs for cybersecurity could help improve future competency assurance programs.

Do high profile heavy sentences on online criminals deter others from crime?

In 2015, the founder of a well-known cryptomarket known as "Silk Road" was sentenced in the United States of America. The judge referenced "deterrence theory" when rendering the sentence, stating that highly publicized punishment may greatly deter crime. We don't know how this theory applies to high-profile convictions in illegal cryptomarkets. Ladergaard collected trade trends information in two cryptomarkets every day for 10 months to test this idea. He then compared the trends with media coverage on cryptomarkets to look at the lag between events, media coverage and changes in market trading. Neither the media coverage of the arrest nor the sentencing of the Silk Road founder appeared to have a deterrent effect on cryptomarket trading. In fact, the effect of the sentencing seems to run contrary to the intention of the court to deter crime, as trade on online markets actually increased with media coverage. It seems questionable to assume that a strong, publicized penalty will deter online crime. It may be more helpful to better understand how online communities function and how decisions about criminal acts take place online.

Can botnet hunting systems work better without our advice?

The methods that make botnets work are constantly changing, which makes them hard to detect. There are systems that can detect botnets through their communications, but longer-term solutions for constantly evolving botnets are needed. Haddadi and Nur Zincir-Heywood compared five botnet detection systems that used different malware detection methods. They tested the five systems on 25 malware collections for malware detection performance. They found two systems that outperformed the others. These two systems utilized machine learning, and were less dependent on predefined botnet malware knowledge than the others. Because malware is constantly evolving, botnet detection systems that depend on predefined features are at a disadvantage. Malware detection that augments predefined knowledge without relying on that knowledge could provide advantages in detecting unknown botnets.

Could an IoT device worm cause problems?

Smart light bulbs are Internet of Things (IoT) devices that can be used to remotely control lighting. IoT networks may be susceptible to similar attacks as network-connected computers. Ronen et al. were able to overcome the security measures of one popular smart light bulb brand. They created an infection that could spread from one smart light bulb to another. The researchers developed their own malicious firmware updates that were accepted as being trustworthy by the light bulbs. This infection could allow an attacker to not only control the lights from a distance, but also be able to use the devices as a part of a bigger network attack. The increasing popularity of IoT devices increases the risk of harm created by poor device security. It is important to keep security in mind when developing IoT technologies.

What should a "Science of Security" be?

There has been a push to make computer research more scientific and develop a "Science of Security". The principle concepts of sciences involve either the observation and measurement of the world or the logical derivation of facts in order to make predictions. Shortcomings in either of these approaches mean that a claim can be either absolutely true or applicable to the real world. In a world where no claims are true it is difficult to tell which claims are sensible. A way to define a sensible claim is to define how to disprove it. In this way, you can at least tell when it is false. When compared to other scientific fields, security research generally falls short when it comes to respecting these core scientific elements. Security tends to confuse types of claims; make unfalsifiable claims, failing to test; make unclear claims and assumptions; and seek to confirm rather than to disprove claims. Security research would benefit from a better distinction between the types of claims made and more rigorous testing of the application of mathematical security proofs to real-world systems. A security science should also be more attentive to unsupported assertions, undocumented assumptions and authority-based arguments, while focusing on refutation with evidence-supported claims.

Are banking apps safe from malicious code injection?

Mobile financial applications can help people with online banking, shopping or money transfers. These applications usually have multiple security measures to protect sensitive data, sometimes including "self-defence mechanisms." Kim et al. wanted to learn more about how these self-defence mechanisms work and how to identify them. The researchers found that self-defence mechanisms have two main methods: checking for "device rooting" and displaying an alert dialogue. Device rooting gives access to the highest level of administrative privileges on the device. The research team built a tool that can locate the self-defence mechanisms in the code. They then selected 200 random apps from the finance category of the South Korean Google Play app store and categorized the apps on whether they checked for device rooting or app integrity. Using the tool they built, the researchers were able to identify and consequently bypass the self-defence mechanisms in nearly 90% of the apps they studied. Current self-defence mechanisms are ineffective on their own and require additional measures to ensure the security of financial application users.

Do botnet controllers prefer certain ISPs for their Command and Control servers?

A botnet is a network of remotely-controlled computers managed by a bot master. The bot master issues instructions to the botnet from a Command and Control (C&C) server. Bot masters may rent or illegitimately gain control of servers from various hosting providers to operate their botnets. Hosting providers have different characteristics, such as price or legal regulations. Tajalizadehkhoob et al. wanted to know whether bot masters seem to prefer certain types of server hosting providers. The researchers gathered information from over 45,000 hosting providers. They discovered that C&C domains seem to be concentrated in a small number of providers. They also looked at a number of trends for different hosting providers, including size and legal regulations as well as popularity, time in business, best price and software vulnerabilities. The researchers found that bot masters seemed to take a fairly random approach when choosing a host though there appeared to be more C&C servers on hosting providers that were larger, more popular, more established, used more vulnerable software and had better prices. Hosting providers with larger IP and domain name space size should be aware they appear to host more C&C servers than others.

How do police manage revenge pornography complaints for high-school age adolescents?

Increased access to social media among youth brings new opportunities for socialization and intimacy, but also for non-consensual intimate image sharing. In Canada, the presence of child pornography laws creates a complex legal landscape for cases of non-consensual intimate image sharing among minors. Dodge and Spencer were interested in police interpretations and responses to non-consensual intimate image sharing among Canadian youth. The researchers interviewed 70 members of sex-crime-related units in police organizations from 2014 to 2016 on their experiences with such cases. Police officers seem to believe that child pornography charges were not designed for offenses occurring between minors and are inappropriately harsh for cases of underage non-consensual intimate image sharing. Many officers seemed to use their discretion to determine an alternative response; such as education. However, some officers voiced ideas that placed blame on the victims of non-consensual intimate image sharing, which may shame and discourage youth from seeking support. There is a need for alternative non-criminalizing responses that educate and refrain from blaming victims to help police organizations cope with underage non-consensual intimate image sharing.

Tracking Phishing Attacks Over Time

Phishing is a large problem for computer security. It is a computer attack that involves the use of messages designed to get people to do things for the benefit of an attacker. Effectively dealing with phishing requires a clear and deep understanding of the messages in order to effectively identify and avoid them.

Cui and a team combining industry and academia worked to better understand this type of attack by focusing on phishing websites. The researchers collected nearly 20,000 websites to study by downloading pages each day for 10 months. Rather than comparing phishing websites to legitimate websites or looking at the characteristics intrinsic to phishing messages, the researchers compared the phishing websites with each other. They wanted to see if they could find common elements in the malicious webpages. The researchers compared the aspects of the tag elements in the webpage code to find structurally similar pages. They then grouped the similar websites together.

The research team compared a total of 19,066 phishing sites. As expected, most of the websites had a short lifetime. Surprisingly, 20% of the studied websites lasted for more than a month. Comparing the content of phishing websites provided useful insights. The changes attackers make to phishing websites to avoid detection are often subtle, such as switching the page to a different domain or subdomain. They are less likely to change the content itself of pages. In fact, the researchers found that 90% of the pages were repeats of an earlier attack.

By classing pages by content, it is possible to see that although instances of phishing change very rapidly, their classes are used over a much longer period. It is likely that recreating the websites or messages is more demanding and costly than creating new instances at varied domains. Detecting phishing based on attack classes could possibly shift the burden of change needed to avoid detection so that it is more costly for attackers.

Detecting phishing messages based on the structure of the content in the pages could provide an additional screening tool. This method could effectively identify a large proportion of phishing pages. It overcomes shortcomings of filtering based on a URL or domain name. Requiring extensive modification of phishing page content increases the difficulty for attackers and could slow down their rate of adaptation to filtering technologies.

Detection methods that focus on page content could shift the burden on criminals to make avoiding detection more costly.

Cui, Q., Jourdan, G., Bachmann, G. V., Couturier, R. & Onut, I. (2017). "Tracking Phishing Attacks Over Time." Proceedings of the 26th International World Wide Web Conference (IW3C2, 2017), Apr 03 - 07, Perth, AU. pp. 667-676.

Follow the traffic: Stopping click fraud by disrupting the value chain

A large volume of the content readily available on the Internet is free. Advertising revenue is important in facilitating free content and for development of the Internet, as it provides a source of income for content creators. The online advertising industry has become a complex arrangement of organisations including advertising networks and syndications. Organisations pay advertisers for promotions based on the number of views and clicks on an advertisement, or for directly resulting sales. Fraudsters create fake traffic which imitates human views or clicks on advertisements, but leads to no real revenue. Advertising fraud is a big business that costs companies an estimated 6 billion dollars annually. This loss threatens the provision of free content on the Internet. We could potentially reduce advertising fraud by disrupting the relationships and connections that enable criminal activity. A better understanding of how these links work would help determine how to most effectively disrupt the illicit flow of money.

Faou et al. studied an advertising click fraud botnet to shed light on how this criminal enterprise works. By using a cross-disciplinary team of researchers, they were able to study the function of a click fraud botnet over a period of seven months to better understand the network and relationships involved in this criminal enterprise. A click-fraud botnet consists of a network of malicious software illegally installed on computers around the world. This malware generates traffic that imitates people clicking on advertiser links. The researchers gathered small amounts of traffic generated by malicious software from a well-known click fraud botnet. From this sample, the team was able to build a picture of the malware-generated advertising traffic. They then matched this data to individual operating entities by grouping the target URLs by similarities in their registry information, passive DNS data, tracking codes and the requested pages. This allowed them to find 225 actors potentially involved in the network. The researchers then mapped the relationships shown by the botnet traffic and measured the direct and indirect influence of actors on the network.

Potential Effects of Insulating Legitimate Operators from Fraud Network

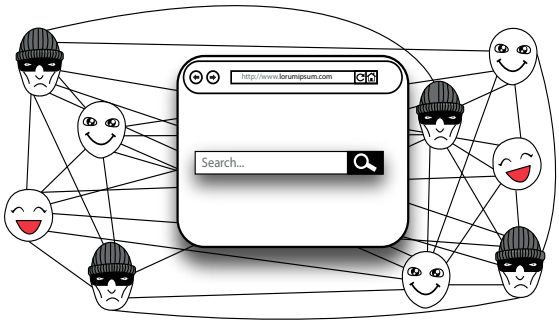

Current Click Fraud Networks

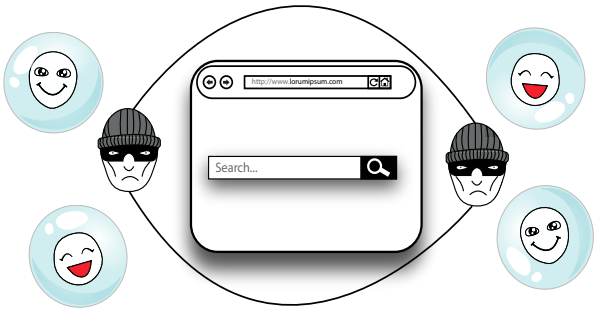

Network After Legitimate Operator Insulation

The key players of the network were identified by looking at the ties between the actors. The research team found which actors they could target to disrupt the network. By removing just four actors from the network, the researchers found they could disrupt 80% of the network. Fragmenting the network makes effective monetization of the network more challenging for fraudsters. However, real world practicalities can make effectively removing the most prominent members difficult. It may be more feasible to focus on the businesses that interact with legitimate customers as well. The livelihood of these companies relies on a level of trust between them and their customers. Removing these companies from the botnet network would make monetization for fraudsters difficult by impacting 50% of the network.

Operations to takedown illegal operators of advertising fraud botnets can be highly effective in the short term. However, encouraging legitimate online advertising businesses to disconnect themselves from fraudulent operators and illegitimate advertising traffic would be a longer-term strategy for botnet network reduction. This may reduce harm to advertisers, their customers and create a healthier online advertising industry.

If legitimate advertising networks isolated the fake click market it could have a major impact on click-fraud.

Faou, M., Lemay, A., Décary-Hétu, D., Calvet, J., Labrèche, F., Jean, M., Dupont, B. & Fernandez, J. M. (2016). "Follow the traffic: stopping click fraud by disrupting the value chain." Proceedings of the 14th Annual Conference on Privacy, Security and Trust (PST, 2016), Dec 12 - 14, Auckland, NZ.

All That Glitters Is Not Gold: On the Effectiveness of Cyber Security Qualifications

The problem of the shortage of skills in cybersecurity is generally considered to be solvable through industry professionalization, competency requirements and training programs. Competency gained through training is often expressed in terms of qualifications. For cybersecurity, the methods used to assess competency vary between qualifications. The importance and variety of assessment in qualifications merits an analysis of the effectiveness of these differing examinations in meeting industry needs.

Knowles et al. reviewed 74 industry-focussed cybersecurity qualifications examinations to see which methods of assessment they use. They found there were five distinct testing methods: "Multiple Choice Examinations"; Narrative Form Examinations"; "Oral Examinations"; "Virtual Lab Examinations" and "Employment History and Qualifications Reviews." The researchers surveyed 153 industry stakeholders in a variety of roles to see how effective they considered these testing methods to be at measuring cybersecurity competence. The survey enquired about perceptions of which of the tests were considered to be effective and which combinations of testing were economical. Multiple Choice Examinations were dominant, with 60 of 74 of the examined qualifications making use of this method. Furthermore, 36 of the qualification programs used Multiple Choice Examinations as the sole form of assessment. Virtual Lab Examinations and Employment History and Qualification Review were the next most common methods, appearing in 21 and 20 of the 74 programs, respectively. However, these methods were much less often used as the sole form of assessment than Multiple Choice Examination.

It is important to note the examinations within each method vary in how intensive and exacting they are. Virtual Lab Examinations ranged from short focused single task oriented assessments to 72 hour intensive examinations. Likewise, Employment History and Qualification Reviews could range from simply requiring a minimum number of years of industry experience to needing specific qualifications or types of experience.

Efficacy perception against industry use

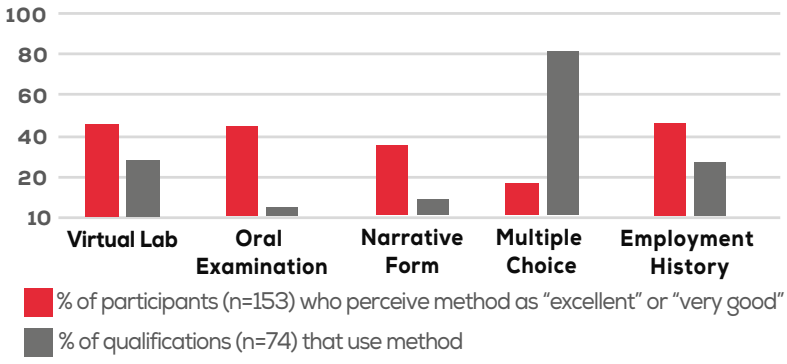

Survey respondents did not seem to perceive Multiple Choice Examinations as very effective. Nearly half (45.5%) of respondents rating perceiving it as a fair or poor competency assessment method. Respondents rated both Employment History and Qualification Reviews and Virtual Lab Examinations more highly with over 45% marking them as "good" or better.

Respondents considered the combination of Oral Examinations and Employment History and Qualification Reviews to be the most cost-effective combination for cybersecurity competency assessment. Each of these methods also often appeared in combinations considered to be cost-effective. The second highest rated assessment groupings of Multiple Choice Examinations together with Employment History and Qualification Reviews combines assessments considered as opposites in efficiency. This was also the most common combination used in practice, with 13 of the 30 qualifications with composite examinations using this combination.

The least effective assessment method appears to be Multiple Choice Examinations; however, this method is required for 81% of qualifications and is the only assessment method in 47% of qualifications. A large proportion of current cyber security industry qualifications use methods perceived as being ineffective. Reappraising the examination methods that industry qualifications use could help in understanding the balance between testing cost and effectiveness. Understanding the appropriateness of testing and the relative costs for cybersecurity could help in underpinning future competency assurance programs.

Current qualification testing appears to be considered neither effective nor cost effective. This should be remedied if industry is expected to close the skills gap.

Knowles, W., Such, J. M., Gouglidis, A., Misra, G. & Rashid, A. (2017). "All That Glitters Is Not Gold: On the Effectiveness of Cyber Security Qualifications." IEEE Computer, 1-10.

We Know Where You Are, What You Are Doing and We Will Catch You: Testing Deterrence Theory in Digital Drug Markets

There are online e-commerce markets that connect thousands of vendors with even larger groups of customers around the world, but intentionally obstruct the tracking of their vendors and customers. These e-markets are referred to as "cryptomarkets." These "cryptomarkets" are used by some to trade in illicit goods and services and provide access to a greater variety of sellers and product than was possible in street trade. In 2015, the founder of a well-known cryptomarket called "The Silk Road" was arrested and convicted in the United States of America. When handing down the sentence, the judge referenced "deterrence theory."This theory states that punishment deters crime and that highly publicised punishment may provide an even greater deterrence. Deterrence theory is based in the notion that people consider the risk of punishment before committing a crime. It follows that a greater perceived risk of getting caught or receiving harsh punishment would cause some to reconsider committing this crime causing the crime rate to drop. Past research on deterrence theory in practice has shown mixed results, suggesting that deterrence has just a weak effect on crime. Cryptomarkets offer a new opportunity to examine this theory. These markets are unique in how decisions about trades are made and how trust is formed and maintained. The Silk Road case presents an opportunity to test ideas about deterrence in the context of cryptomarkets.

Ladegaard wanted to test two ideas: if media coverage of advances in the investigation abilities of law enforcement agencies would reduce crime and if the high-profile conviction and sentencing of a key person would reduce trade in similar illegal markets. To do so, he collected mandatory feedback on trades in two cryptomarkets every day for 10 months. This data represents the trade trends for those markets. The researcher then compared these trends with reports collected from media sources and relevant forums to understand the lag between events, media coverage and changes in market trading. Ladegaard also looked into the commentary on online forums relating to cryptomarkets to provide further insight.

Neither the media coverage of the arrest nor the sentencing of the Silk Road founder appeared to have a deterrent effect on cryptomarket trading. Trade on online markets actually increased with the media coverage. However, this is not to say conclusively that the verdict and sentencing were completely ineffective as deterrents. We cannot know if there would have been even more trade if the trial result was an acquittal. Perhaps the sentencing did deter existing traders, but this effect was outweighed by an influx of new traders. Nonetheless, there does seem to be an effect from a high profile sentencing that runs contrary to the intention of the court to deter crime.

The forum commentary shows the prosecuted founder of Silk Road in different ways. Often, commentary portrayed him as an idealistic pioneer martyred for the cause of online freedom or for being an advocate against the war on drugs. Other comments depicted the founder as a poor child caught out of his depth for not knowing the basics of avoiding apprehension. These types of comments suggest a community that sees illicit online trade as a necessary expression of a higher moral standard than the current law, allowing their justification of illegal activity. Alternatively, the idea that human error led to the capture and conviction of the founder somewhat restores faith in the Cryptomarket mechanisms that protect traders.

The assumption that a strong penalty will deter further crime online is questionable. In this case, it seems that illicit trade increased after media coverage of the arrest and sentencing of the Silk Road founder. Decisions based on a better understanding of how online communities function and how decisions about criminal acts take place online could provide better outcomes.

The deterrence effect of high profile cases and harsh penalties should be better understood before provided as justification for sentencing.

Ladegaard, I. (2017). "We Know Where You Are, What You Are Doing and We Will Catch You: Testing Deterrence Theory in Digital Drug Markets." British Journal of Criminology.

Botnet behaviour analysis: How would a data analytics-based system with minimum a priori information perform?

Botnets are an evolving problem. A botnet is a network of computers that have been infected with malware allowing them to be controlled without their owners' permission. The systems and methods that botnets use to operate change continuously; making them hard to detect. However, for practical purposes they must have an automated updating process. This provides a potential method for detecting botnets through their communications. Innovation in botnet operation make it difficult to identify new botnets based on knowledge of past botnets. Developing long-term solutions requires an understanding of what makes a botnet detector effective.

Haddadi and Zincir-Heywood compared five botnet detection systems. The systems chosen used different methods to detect malware. They used regulaly updated sets of rules, analysed message content, or analysed content extracted from the flow of communications traffic. Four of the detection methods were based on knowledge of existing malware.

The researchers came up with comparative tests for five systems: Snort, BotHunter, Trananlyzer-2, FlowAF and a packet payload-based system.

Snort	Intrusion detection and prevention system that matches data packets to predefined signatures (rule sets) based on a priori knowledge. Publicly available.
BotHunter	Based on the idea that all botnet infection processes are similar and so it detects packets related specific bot actions at different stages of its life to better detect infected machines. Publicly available.
FlowAF	Uses machine learning algorithms to classify traffic as originating from botnets traffic based on the time between packets, or the flow intervals.
Trananalyser-2	Uses a machine learning algorithm to identify malware based on the features of the flow of packets, rather than analyzing individual packets. The system chooses features that provide the most insight itself rather than using ones selected by an expert.
Packet payload system	Attempts to classify as matching known types of malware packets based on their features, such as the port, protocol, or size of the packets.

Snort, BotHunter, the packet payload–based system, and FlowAF are systems that use expert knowledge of malware to define the rules and features for detection. The Tranalyzer-2 flow-based system used a minimum of established knowledge to extract a wide range of defining features for malware. The researchers tested the 5 systems by applying them on 25 publicly available malware collections. The results suggested that the Tranalyzer-2 flow-based and FlowAF systems outperformed the other systems. Using specific features to detect malware could limit how well a system works when faced with new threats. It makes sense that it should be able to change how it detects malware to suit what is happening in its environment. Classification methods that can detect unusual traffic behaviour without having to rely on predefined knowledge would be better at adapting over time. It seems that most useful features for malware detection are those related to the communications traffic flow. More particularly, the space between the arrival of data packets appears to be important. This implies that the flow of traffic in botnets, even decentralized botnets, is different enough from normal user behaviour to be detectable.

Relying on predefined features of botnets or malware places detection software at a disadvantage. Malware changes and evolves constantly, meaning that past experience has limited benefit in new scenarios. Malware detection that augments predefined knowledge but not relying on that knowledge can assist it to detect unknown botnets.

Malware discovering systems augmented with detection that is based on discovery rather than expert opinion could be more effective in detecting unknown botnets.

Haddadi, F. & Zincir-Heywood A. N. (2017). "Botnet behaviour analysis: How would a data analytics-based system with minimum a prioria information perform?" International Journal of Network Management.

IoT Goes Nuclear: Creating a ZigBee Chain Reaction

Smart devices and Internet connected devices have increased greatly in popularity creating an Internet of things (IoT). One of these devices is the smart light bulb, which is used to remotely control lighting. They have been sold in great numbers and installed in a large number of homes in modern dense cities. Smart devices and their networks may be susceptible to the same types of issues as network connected computers. Given the density of modern cities and the potential range of IoT communication networks, a specific attack could possibly hop across networks to infect a large area.

Ronen et al. tested the possibility of smart light bulbs being vulnerable to a self-propagating malware attack, or a worm. They developed and tested a method for overcoming the wireless network security standard used by the Phillips Hue smart light system, called Zigbee. The Phillips Hue lighting system has a number of security systems to defend against attacks. Primarily, it uses encrypted communications and a proximity check mechanism to prevent devices outside the home from performing dangerous operations. However, the researchers were able to overcome the encryption by using a side channel attack based on an analysis of the power used in the device hardware. This attack measures the use of power through the device as the encryption/decryption operation is happening to learn more about when processing is being performed. This provides insight into how many and what type of functions are likely being performed and eventually allowed the researchers to find a feasible method of determining the encryption keys.

A bug in the code of the lights allows the device to be reset without passing the proximity test. The reset device will then connect to another network. This means that a light can be infected from another light allowing the worm to hop across networks it would have otherwise been too far away from.

The researchers created their own malicious firmware updates for the devices that were accepted as being trustworthy by the device. They also showed that it would be possible to infect a device from a distance using a more powerful transmitter or from a drone platform. To reduce the risk of a worm escaping the test environment and causing damage, the researchers did not combine all of the exploits into a fully- functional worm.

This research shows the importance of security in the development of IoT devices, their protocols and the certification processes for those protocols. Using unique encryption keys for each bulb would prevent a simple worm from spreading. These types of devices might seem low risk; however, attackers could use them to jam other devices on similar networks, transport sensitive data out of networks or potentially harm people. The increasing ubiquity of IoT increases the potential harm they could cause. Consequently, the risk tradeoff between convenience and security should be rebalanced to consider the greater harm.

IoT communication networks are susceptible to malicious exploitation, possibly providing attack vectors to or data exfiltration from other more sensitive networks.

Ronen E., O'Flynn, C., Shamir, A. & Weingarten, A-O. (2017). "IoT Goes Nuclear: Creating a ZigBee Chain Reaction." Department of Computer Science and Applied Mathematics: Weizmann Institute of Sciences Rehovot, Israel.

SoK: Science, Security, and the Elusive Goal of Security as a Scientific Pursuit

Measuring and analyzing security in a way that reliably shows progress is a unique and difficult challenge. Computer security research is unclear. This lack of clarity has led to a movement to put more science into security; or develop a "Science of Security." However, exactly what scientific security research would, could or should look like is vague. Herley and van Oorschot offer some insight into this problem by highlighting aspects from the history and philosophy of science that are relevant to security research. They discuss the current science of security and how security research has not adopted lessons learnt from other sciences.

Inductive and deductive statements are different types of knowledge claims. Inductive statements are claims based on the observation of a real-world event and are inherently real. However, the claims are only true for those things observed and does not provide insight into things not yet observed. Alternatively, deductive claims are based on a self-evident truth and are absolute. The nature of a claim based on a preconceived premise is that it is abstract from reality. This means that it is not as applicable to the real world and is not able to explain anything new. Either kind of statement is compromised by being either true only for things observed or under certain conditions. There are no statements about the world that we can be absolutely certain are true. If no statements are certainly true then all statements could be considered as being equal, regardless of how sensible they are.

One way to overcome this problem, and make sensible statements, is to make the statement falsifiable. In other words, such a statement must be able to be proven to be wrong. This is what differentiates a scientific claim. Consequently, claims about things that cannot be observed or objectively tested are unscientific. Statements about things not physical such as religion, metaphysics or mathematics do not meet this criterion. Deductive methods that describe the world, such as geometry, may parallel the physical but are still separated concepts. The alignment between concept and reality must be tested to be confident about how it lines up in different situations. Scientific inductive statements are those that can be disproved. Scientific deductive statements are those that have an alignment to reality that can be disproved. Science is then able to be constructive, as something can be known and predictions about the unknown can be made. This understanding is consistent across much of the scientific world and is the basis of the method of: hypothesize, predict, validate.

Computers have evolved from a well understood and defined environment to a mess of interactions, unstructured data and adversarial input. It is now a chaotic environment not unlike the real world. Computer security has developed a great deal in the past four decades, but its scientific aspirations have suffered from disparate approaches. Security is often reliant on strict models or mathematical proofs that can fail when presented to the real world. When compared with other scientific fields, security research generally falls short by confusing inductive and deductive claims; making unfalsifiable claims, not testing; not making claims and assumptions clear; and seeking to confirm rather than to disprove claims. The security research community would benefit from a focus on core scientific elements, a greater commitment to clearly distinguishing between inductive and deductive statements to reduce confusion and rigorously testing the coupling between mathematical proofs of security and real-world systems.

Researchers should use evidence of effectiveness to justify security measures and avoid unfalsifiable claims or circular arguments. Properties of computing and its security may be unique but not so much as to exempt it from scientific approaches. A scientific approach does not mean that security should focus on laws or proofs like physics or cryptography. A prejudicial approach to developing security science could result in using inappropriate methods resulting in precisely measured, mathematically proven falsehoods. A harmonized application of both theory and measurement to security are needed for progress on the diverse set of problems in security research. This science should be attentive to unsupported assertions, undocumented assumptions and authority-based arguments while prioritizing efforts at refutation with evidence-supported statements. Further, placing research in context reduces the risk of providing orphan security components that never fit into a complete security solution.The security community is not learning from history lessons well-known in other sciences. Simply wishing for a "Science of Security" will not make it happen. Security researchers need to learn and adopt more scientific methodologies.

A "Science of Security" will come from collectively understanding and applying core scientific principles to making and testing claims about security issues.

Herley, C. & van Oorschot, P.C. (2017). "SoK: Science, Security, and the Elusive Goal of Security as a Scientific Pursuit." Proceedings of the 38th IEEE Symposium on Security and Privacy (SSP, 2017), May 22-24, San Jose, CA.

Breaking Ad-hoc Runtime Integrity Protection Mechanisms in Android Financial Apps

There are many different financial applications available for Android devices across the world. These mobile applications, or apps, provide people with online banking, money transfer, shopping and other financial services. Financial apps usually have multiple security measures in place to protect sensitive data. These measures sometimes include "self-defence mechanisms." These mechanisms check for a compromised device or application platform tampering. They stop the app from loading and provide a warning message if something seems suspicious. However, little is known about how these mechanisms work. Android financial apps often obscure their code, making observation difficult.

Kim et al. wanted to know more about self-defence mechanisms and the apps that use them. More specifically, they looked at how these mechanisms determine if an application platform or device is compromised and how to precisely identify these mechanisms.

They found two main steps of self-defence mechanisms: checking for "device rooting" and displaying an alert dialogue. Device rooting is gaining access to the highest level of administrative privileges on the device. The researchers developed a tool that generates a human-readable call graph of the methods related to self-defence mechanisms. This tool locates the self-defence mechanisms in the code of apps.

The researchers then selected 200 random apps from the Finance category of the South Korean Google Play app store. To see which apps used self-defence mechanisms, the researchers tested each app in two conditions: on a "rooted" device; and after altering the app's binary code. If an app failed to launch or execute properly, the researchers concluded that the app used a self-defence mechanism. They then classified apps on whether they checked for device rooting or app integrity.

Number of apps bypassed

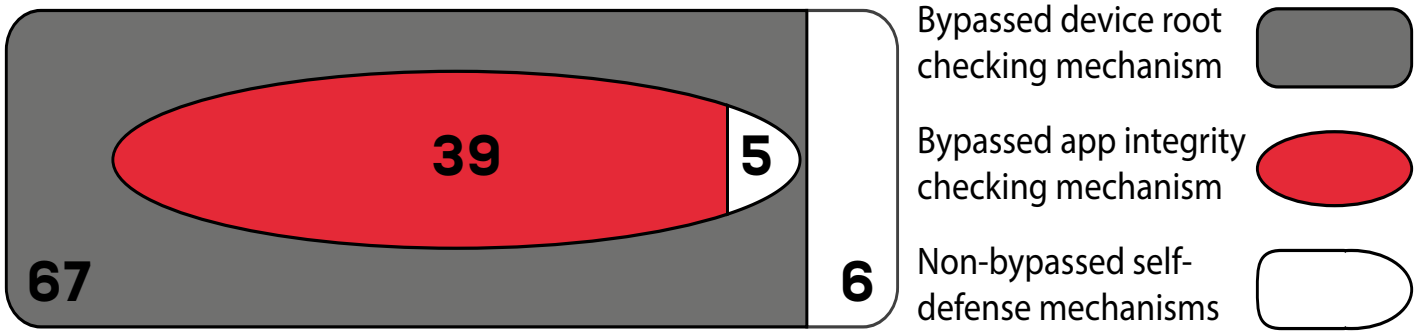

The research team found 76 apps in total that used self-defence mechanisms: 73 of which performed device rooting checks and 44 that checked for app integrity. The researchers then ran their tool on the 76 apps and found that they could successfully bypass the self-defence mechanisms of 67 out of 73 apps utilizing device rooting checks and 39 out of 44 of the apps checking for app integrity. The researchers have indicated that it is possible to bypass nearly 90% of the observed self-defence mechanisms.

The researchers found a way to identify and consequently bypass the self-defence mechanisms in nearly 90% of the apps they studied. This suggests that current self-defence mechanisms are ineffective on their own. There is a need for additional security mechanisms to ensure application and platform integrity for Android financial apps with self-defence mechanisms.

Current mobile financial application defences against malicious code insertion appear to be inadequate.

Kim, T., Ha, H., Choi, S., Jung, J. & Chun, B-G. (2017). "Breaking Ad-hoc Runtime Integrity Protection Mechanisms in Android Financial Apps." Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security, April 02 - 06, Abu Dhabi, UAE.

The Role of Hosting Providers in Fighting Command and Control Infrastructure of Financial Malware

A botnet is a network of remotely-controlled computers managed by an operator, known as a bot master. Bot masters use a Command and Control (C&C) infrastructure to control their botnets. The bot master gains control of target hosts by infecting them with malware. The infected machines connect to a "server" so that the bot master can issue instructions. Bot masters can rent or illegitimately use servers to run their C&C infrastructure. Many hosting providers have efforts set in place to prevent people using their service for C&C infrustructure. Little is known about what motivates bot masters to make decisions when choosing C&C servers, such as price or risk of discovery.

Tajalizadehkhoob et al. wanted to know whether bot masters seem to prefer certain types of server hosting providers or whether the distribution of C&Cs across the industry is random. They looked at trends in different hosting locations of C&C domains for 26 different malware families known to be involved in attacks on financial services.

The researchers started with a database of C&C domains and a database of hosting providers. Together, these databases provided the domain name and IP information of C&C domains in 109 countries from 2009 to 2016. The researchers associated the known IP addresses to their respective hosting providers. In total, they located and identified over 45,000 hosting providers. They then analysed the different trends for the hosting providers. The researchers looked at the size and legal regulations of every hosting provider. They then looked at the popularity, longevity, best price and software vulnerabilities of smaller samples of the hosting providers. The researchers also reviewed the "uptime" of C&C domains by measuring the number of days between the first and last observation of the domain in their dataset.

The results show a general increase in hosting providers over time. The researchers found just 30% of providers hosted 80% of the C&C servers. Most of the domains in the top 20 were located in the USA and Western Europe. The majority of C&C domains seem to be quite short-lived, but some domains do appear to stay hosted for longer periods of time; such as over a year. Bot masters seemed to have little or no preference for hosting providers that leave C&C domains up longer. However, there was an increased presence of C&C servers on hosting providers that were larger, more popular, more established, used more vulnerable software and had better prices. However, all things considered, the researchers concluded that bot masters choices for C&C hosting seems to be random, with the most probable deciding factor being the size of the hosting provider.

It appears that bot masters take a fairly random approach when choosing hosting providers for their C&C servers. Though some may take into consideration a hosting provider's popularity, longevity and pricing, it seems to be the size of the host that matters the most. Hosting providers with larger IP and domain name space size should be aware that their size seems to increase the chance of their business hosting more C&C servers than others.

Bot masters don't appear to have preferred service providers for their C&C servers.

Tajalizadehkhoob, S., Gañán, C., Noroozian, A. & van Eeten, M. (2017). "The Role of Hosting Providers in Fighting Command and Control Infrastructure of Financial Malware." Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security, April 02 - 06, Abu Dhabi, UAE.

Online Sexual Violence, Child Pornography or Something Else Entirely? Police Responses to Non-Consensual Intimate Image Sharing among Youth

The rise of social media and an ever-connected lifestyle offers new opportunities for socialization and intimacy. Unfortunately, these opportunities also present new ways for youth to engage in online sexual violence. One form of online sexual violence that has been getting a lot of attention is non-consensual intimate image sharing. In Canada, the distribution of an intimate image of a person without consent is a criminal offence. The distribution of intimate images of a person under the age of 18 is also an offense, as it constitutes child pornography. This can create delicate situations for police officers responding to complaints of non-consensual distribution where the persons involved are minors. There are multiple cases of Canadian courts charging underage non-consensual distributors of intimate images of other minors with child pornography offenses. However, little is known of the experiences of police officers who must respond to such issues.

Dodge and Spencer were interested in police interpretations and responses to non-consensual intimate image sharing among Canadian youth. They conducted 70 individual interviews and two focus groups with members of sex-crime-related units in police organizations from 2014 to 2016. The researchers selected participants from police organizations in a way that reflected the varied urban landscapes of Canada. The interview questions covered various open-ended topics, such as: their experiences with investigating sex crimes, the use of technology, and their perceptions of the criminal justice system in response to sex crimes.

Police officers appear to find child pornography charges inappropriately harsh for cases of underage non-consensual intimate image sharing. According to many participants, child pornography charges were not designed for offenses occurring between minors. Some officers also mentioned that intimate image sharing, whether consensual or non-consensual, is commonplace and normalised among youth. Therefore, child pornography charges would overly criminalise far too many youths.

The majority of officers interviewed did not mention the non-consensual intimate image sharing law as an alternative to a child pornography charge. According to some participants, a victim's parents or school personnel may push for the use of criminal charges. In contrast, some officers describe the reaction of the victims themselves as not wanting to criminalize their peers.

Most officers seemed to prefer using their discretion to determine what they consider to be the best possible alternative response. Forms of police officer discretion varied from using scare-tactics to instill fear in youth, to educating youth about the law or providing school officials with educational material and information. While education may be a positive alternative response, certain officers voiced ideas surrounding online sexual violence that were problematic. Some participants expressed views that effectively placed blame on the victims of non-consensual intimate image sharing. These officers seemed more focused on changing the behavior of victims who consensually created these intimate images instead of the behavior of the perpetrators who non-consensually shared the images. These responses may shame youth for creating intimate content and discourage victims from seeking support.

There is a need to better define the difference between consensual and non-consensual intimate image sharing among youth. Proactive educational approaches that recognize the importance of privacy and sexual consent and refrain from blaming victims of non-consensual intimate image sharing should be privileged. These types of alternative, non-criminal responses may also help officers with the additional legal burden posed by an increase in online sexual violence among youth.

There seems to be a space where laws for non-consensual intimate image sharing and child pornography overlap leaving police officers exercising discretion without clear guidance.

Dodge, A. & Spencer, D. C. (2017). "Online Sexual Violence, Child Pornography or Something Else Entirely? Police Responses to Non-Consensual Intimate Image Sharing among Youth." Social & Legal Studies, 1-22. doi:10.1177/0964663917724866

serene-risc.ca

Smart Cybersecurity Network (SERENE-RISC) is a knowledge mobilization network that aims to be a cybersecurity exchange that is recognized as an open, inclusive, and unbiased forum to unlock the value of knowledge on cybersecurity risks, threats and solution for all Canadians. We are a network of networks that provides publications, tools, events, professional development and education. We are always looking for partners to collaborate to make cybersecurity better, so please contact us to find out more.

konnect.serene-risc.ca

We are building the Canadian source for summaries of research, opinion pieces, video presentations, fraud bulletins, public awareness materials and more. You can filter the collection by type of content, click on keywords or search. The website provides lots of original conent produced by us and in collaboration with partners. To save your time, ee are sending an email regularly a summary of the new content to save you time and make it easier to find content. You can subscribe to this list online at: **http://konnect.serene-risc.ca/subscribe-abonnement/**

We will be looking for contributions to this page from our community, so if you have an idea for a piece that you would like to share please let us know.

You can help by:

- Subscribing to the regular update online at: **http://konnect.serene-risc.ca/subscribe-abonnement**
- Following **@SERENE_RISC** on twitter and retweeting,
- Joining the Linkedin Group and submitting or commenting on posts, and
- Posting links to **konnect.serene-risc.ca** content on other platforms you are involved with.

cybersec101.ca

Want to provide basic cybersecurity classes but don't know how or have the time to get started?

We have developed evidence-based materials that can help you quickly put on training programs. Find resource sheets, presentation videos, class tools and slides to build basic cybersecurity education and awareness sessions at. There are ten modules in French and English with materials to help from basic concepts to practical step-by-steps for better security online.

donate

We are now accepting donations to help us provide more services that are open, accessible, inclusive and ubiased. You can download a donation form at https://www.serene-risc.ca/donation or contact us for more information at info@serene-risc.ca

 @SERENE_RISC  /serenerisc  /serene-risc

The SERENE-RISC Cybersecurity Knowledge Digest

2017 Autumn

Editor-in-Chief: Michael Joyce
Scientific Editor: Benoît Dupont
Editor: Shannon McPhail

Government of Canada Networks of Centres of Excellence | Gouvernement du Canada Réseaux de centres d'excellence

serene risc

Université de Montréal