serene risc Digest Summer 2015

Cutting Edge Research Summaries for Policy-Makers and Practitioners

Are remotely operated surgery robots cyber secure?

Remotely operated robots have an increasing number of applications, among these providing surgical care in remote locations and extreme conditions. Access to these medical services can bring enormous benefit, but is also vulnerable to disruption. In this study researchers tested a surgery robot to evaluate how cyber attacks might impact telemedicine. Attacks made it harder to get the job done and the assumption is that the same challenges would apply to experienced surgical teams in real-world settings. Existing security solutions are feasible to address some of these challenges, but any future design decisions for teleoperated robotics must reconcile the intended benefit of any added function with the potential harm from unintended use.

Can anyone just guess my security questions?

When a user forgets their login information, there must be some way to verify their identity and reset the account. For this verification, some platforms rely on a set of personal knowledge questions defined by the user when they open the account. This study demonstrates that because some questions have common answers, personal questions are not immune to guessing. They are also not immune to forgetting; personal questions provide less than 80% successful recovery. Other ways – such as SMS and email-based account recovery – show more promising results, suggesting that personal questions may not provide the optimal balance of security and memorability.

Are all password strength meters created equal?

Password-strength meters respond to the relative strength of the user`s chosen password, to show how that password will stand up to a hacking attempt. This study tested the password-strength meters on 11 websites. Each site used a different meter and most rely on simple rules. Few detect patterns or predictable changes and the same password choice is often rated differently across the various platforms. These inconsistencies might cause confusion. A user creating a password could be misled, either to believe their weak choice is a strong one or that an otherwise strong password is judged as inappropriate. Users who do not understand the rules or how they are applied might be less cooperative. One part of the solution is for users to understand what makes a secure password and apply those standards in their own choices. Thorough password-strength meters will go beyond character complexity and length rules to consider patterns and common passwords from a wider range of sources.

Is there a way to design systems that support better security decisions?

Designing strong tools is an important step in providing security. A system designed with the user in mind can support an engaged user to make informed decisions about security. Computing systems can leverage the human relationship concepts of trust and comfort to permit users to make their own security decisions. With the Device Comfort approach a device compiles evidence about the degree of comfort with the environment, the user and their task and then presents this information to the user, to advise, encourage and even warn the user about potential actions, in a way that helps the user understand the security implications. Rather than block or override the user instruction, the device asks the user to pay attention before completing the action. By bringing trust to the foreground, the Device Comfort model reminds the user they are making a security decision and supports an informed choice.

How bad is the situation with ransomware?

Recent years have seen an increase in ransomware, malware that locks the affected computer and demands payment from the victim in order to restore their files. This study examined characteristics of the most common types of ransomware observed in the real world between 2006 and 2014. Most ransomware are not very sophisticated and even lack the ability to execute a complete attack; they rely on superficial approaches to disable user access, such as by locking the desktop, rather than more thorough methods like encrypting or deleting files. Several different classes of ransomware cause similar changes in file system activity, with activity that is distinct from non-malware processes. These patterns of activity could be leveraged for early detection and earlier intervention to recover deleted files.

Can we prevent Security Analyst Burnout?

Many organizations dedicate significant resources to security monitoring. The ability to identify threats to security depends on effective analysts so it's troubling that this group is often characterized by high rates of burnout. Sundaramurthy et al. worked with the security analysts from one Security Operation Center (SOC) to identify the factors behind security analyst burnout.
The factors related to Human capital must be carefully managed, as they reinforce each other in a cycle that can be either virtuous or vicious in its effect on analysts. Analyst Skills, Empowerment, Creativity and Growth make up the elements of this cycle. Communication, automation and performance metrics can be powerful tools in managing this cycle where used effectively. Operations can watch for opportunities to contribute to a virtuous cycle of Security Analyst improvement.

Can we automatically find bad code?

A large number of security and forensic applications, especially in a cloud environment, rely on the fingerprinting of executable code. Ahmed et al. present primary research on a technique to identify a known piece of code that is running in an arbitrary location. The authors recognized that previous fingerprinting techniques for Microsoft Windows were not able to work with partial information and required an expert analyst to drive the discovery process. To overcome this problem they proposed a new solution using relocation tables (a directory of where each executable code is located) as the key identifying characteristic. Identifying executable code is an important aspect of security monitoring and forensic analysis. The research provides another tool for code identification, which will become increasingly important in complex environments and investigations. This research could provide enhanced abilities to providers by offering a robust method for automated fingerprinting.

Under what criteria can a cyberattack be ethically justified?

The relatively new phenomenon of attacks and warfare in cyberspace make it difficult to understand the effects of these actions. While international law has established some rules and regulations for cyber attacks, it may be that traditional moral frameworks still apply in cyberspace. Barrett's research builds on the existing legal framework for military cyber operations to include an ethicist's perspective on the usefulness of including both law and the traditional just war doctrine to provide ethical guidance.
The just war criteria stipulates that only a proper authority can respond to a threat against life and life-sustaining properties in self-defence. Difficulties with attack attribution in cyberspace create a challenge in identifying whom to respond to. To be deemed a justified defensive response, the attack must be attributed with absolute certainty,and must be able to discriminate the target, which is difficult in this highly unpredictable environment. International law is insufficient in guiding cyber conflict, however where combined with a moral agent operating within the just war criteria are capable of providing the ethical guidance needed to legitimately conduct military cyber operations.

Why are some security warnings more effective than others?

Modern browsers warn users when they are potentially clicking from an encrypted connection, using the Secure Socket Layer (SSL) to an unprotected connection. This can indicate a potential threat to the user.
Despite the difficulty in assessing the risk highlighted by the warning, many users 'click-through' this message. Felt et al. set out to understand why Firefox has a 37% lower 'click-through' rate on SSL warnings than Google Chrome. In a test where users had to click through a warning twice, the second did little with almost all participants (98%) clicking through both. The use of corporate style or branding had little effect and the use of a warning. This research contributes to the ability of SSL warning designers to create effective warnings that avoid the use of technical jargon by de-emphasizing technical details, include information on ways to mitigate risk, and include a clear default choice.

Are most home broadband routers secure?

Digital Subscriber Line (DSL) routers have become an essential part of home networks as they provide the gateway to broadband Internet.
Niemietz and Schwenk investigated ten different DSL routers that use a Web browser to manage their administrative settings. They used the techniques of Cross-Site Scripting, Cross Site Request Forgery and User Interface redressing to manipulate the victim to click on a malicious link or unknowingly share their user and password information. All of the DSL routers tested were vulnerable to most or all of these attacks. The researchers were able to produce a list of simple countermeasures for manufacturers to use in enhancing the security of their product. The implications of this study are not limited to DSL routers as many other devices share the web interface Internet connectivity feature and possible vulnerabilities. This study provides a necessary security evaluation for manufacturers and purchasers of DSL routers.

To Make a Robot Secure: An experimental analysis of cyber security threats against teleoperated surgical robotics

Remotely operated robots have an increasing number of applications, among these providing surgical care in remote locations and extreme conditions. Access to these medical services can bring enormous benefit, but must also be considered as being vulnerable to disruption. Robots that rely on public converged networks can be targeted in cyber attacks that could pose a risk of numerous harm to: a patient`s health and privacy; the legal responsibility of the surgeon; equipment; and even public confidence in telemedicine.

Bonaci et al. used the Raven II surgery robot to evaluate the impact of possible cyber attacks on telemedicine. They put student 'surgeons' at the robot control console to conduct a simple exercise that is used to train and test surgeons. The attack observed and modified the communication between the control console and the robot.

Three types of attacks are particularly relevant to teleoperated robots:

- Intention modification – The attacker changes the message sent from operator to the robot, with noticeable results (e.g. unexpected or unusual movements or delays);
- Intention manipulation – The attacker changes the feedback message sent from the robot back to the operator, with less obvious indicators; and
- Hijacking – The attacker overrides operator instructions and the robot engages some other action.

Under experimental conditions, all types of attacks made it harder to get the job done. Robot movements were delayed and less fluid, or more 'jerky', and, when under attack, the robots were more likely to perform movements with errors. When the communication was hijacked, the operator lost complete control of the robot. It is reasonable to assume that the same challenges would apply to experienced surgical teams when the system is under a cyber attack.

Existing security solutions could feasibly address some of these challenges:

- Robots could accept instructions only from legitimate command sources
- Data streams could be encrypted and then authenticated through another communication medium (e.g. text, or phone call). Based on initial testing, this encryption requires no noticeable increase in CPU usage, but some increase in memory usage
- The instruction processing rate could be limited to prevent a machine acting on a flood of commands in a short time
- The existing communication standard for surgical teleoperation, the Interoperable Telesurgery Protocol, could be updated with additional security specifications
- Link and network status could be monitored to identify and raise alarms about multiple streams of data or large numbers of out of order packets

Some challenges are not so easily addressed. Future design decisions for teleoperated robotics must reconcile the intended benefit of any added function with the potential harm from its intended use. Consider, for example, the emergency stop feature of surgical systems. This mechanism is important for safety during normal operating conditions, but it is highly vulnerable to disruption and possibly dangerous if hijacked. These tradeoffs should be evaluated carefully.

Surgery is a powerful application for teleoperated robotics, however the potential harms from cyber attack apply to all remotely operated technology. Existing security mechanisms could prevent attacks on remotely operated systems, but must be carefully studied for the unique circumstances of use. Overall those who manage or teleoperate robotic systems must recognize the potential harm of the tools within their control, and be capable of weighing the costs against the benefits.

Telesurgery robots present a real security challenge that can be and should be addressed.

Bonaci, T., Herron, J., Yusuf, T., Yan, J., Kohno, T., & Chizeck, H. J. (2015)"To Make a Robot Secure: An Experimental Analysis of Cyber Security Threats Against Teleoperated Surgical Robots",arXiv:1504.04339.

Secrets, Lies, and Account Recovery: Lessons from the Use of Personal Knowledge Questions at Google

When a user forgets their login information, there must be some way to verify their identity and reset their account. For this verification, some platforms rely on a set of personal knowledge questions such as asking one's father's middle name, favourite food, place of birth etc. The questions and answers are defined by the user when they open the account. To regain access to the account, the user must provide answers that match the responses given initially. Although the questions are by nature personal, the answers selected by the user are neither unique to the individual nor immune to guessing. Some questions have common answers, which increases the likelihood of success of mass guessing attacks.

Bonneau et al. examined how security questions are used in Google accounts, to investigate the tradeoff between security and memorability. To study security, the distribution of `hundreds of millions` of actual user secret answers was used to compare the range of user responses to the real-life distribution of accurate answers. To study memorability, data on 11 million account recovery claims from 2013 was used to evaluate the success of personal questions in account recovery. An important aspect of this particular research is methodological; this study had the benefit of data from Google that shows the actual answers given by real-world users, from which the researchers could see which responses recurred. Further, by comparing their real-world data with crowd-sourced information this study has shown that crowdsourcing gives a reasonably accurate approximation of the distribution of actual answers. This means it is possible to crowd-source the most statistically likely answers to personal questions.

The security of personal knowledge questions concerns the frequency of the most commonly provided answers. Some responses are naturally common, especially among some linguistic groups or in a given country. For example, some names are highly popular among Spanish-speakers, perhaps even more so than the most popular names among Anglophones. Similarly, birth city may be heavily skewed in highly urbanized countries with a few large centres; 39% of Korean speakers claim the same city of birth. In other cases, answers are common because users have provided fake answers. Consider responses to the prompt "first phone number". The real-life distribution of phone numbers dictates no two identical answers. However, almost 3% of the Arabic-speaking respondents provided the same phone number. As this example illustrates, providing false responses actually decreases the security of the personal question; with a single guess, an attacker could successfully hack that 3% of Arabic accounts.

Low memorability means low reliability for successful account recovery. The personal knowledge questions are not a useful tool if the user cannot remember the answers. At best, personal questions provide less than 80% successful recovery for all questions; the rate of successful recovery decreases over time, so the more time elapsed since the questions were set the less likely the user is to remember. Other ways show more promising results. SMS and email-based account recovery increase successful recovery over the personal question method by 20% and 14.5% respectively.

Balancing phrase memorability and security is an open question. Given the limited success with personal knowledge questions, other methods should be explored for more efficient account recovery.

Personal security questions for authentication have a large number of failings and other methods should be explored for more efficient account recovery.

Bonneau, J., Bursztein, E., Caron, I., Jackson, R., & Williamson, M. (2015). "Secrets, Lies, and Account Recovery: Lessons from the Use of Personal Knowledge Questions at Google." In Proceedings of the 24th International Conference on World Wide Web

From Very Weak to Very Strong: Analyzing Password-Strength Meters

A password-strength meter judges how a user-chosen password will stand up to a hacking attempt. On many websites password-strength meters show the user a visual or verbal cue for the relative strength of their password. The intention is to lead the user to select a more secure password. Few password-strength meters provide any explanation of how they evaluate passwords, but common criteria include length, complexity, use of common words or easily typed sequences, and matches to user personal information.

de Carnavalet and Mannan tested 4 million common passwords using the meters from 11 different popular websites covering a range of financial, email, cloud storage and messaging services. The common passwords used were taken from dictionaries leaked from hacked websites and other lists of frequently used passwords, and are also the ones attackers are likely to test.

Each of the 11 websites uses a different meter, most of which rely on simple rules. Few of the password-strength meters included an algorithm to detect patterns and minor or predictable changes. None of the meters used the same set of labels to denote the strength of the password. Further, some of the programs did not enforce strength requirements by rejecting poor passwords, but only informed the user of the relative security of their choice. Because of the variation across meters, the same choice can be rated differently depending on the platform. For example, the choice of `password$1` is rated as follows in each of the noted systems:

Service	Strength	Service	Strength	Service	Strength
Apple	Moderate	Microsoft (v1)	Strong	Google	Fair
Dropbox	Very weak	Microsoft (v2)	Medium	Twitter	Perfect
Drupal	Strong	Microsoft (v3)	Medium	Yahoo!	Very Strong
eBay	Medium	Paypal	Weak	Skype	Poor
FedEx	Very weak				

There are also inconsistencies in the results from a single meter. That is, the rating assigned sometimes changes when the same password is attempted at different times, or the rating might change significantly for minor revisions to the password that may not actually reflect an improvement to the security of the choice.

Inconsistent results might cause confusion. A user creating a password could be misled by the meter's response to their password choices – either to believe their weak choice is a strong one or that an otherwise strong password is judged as inappropriate. If users do not understand the rules or how they are applied, they might be less willing to comply with the guidelines, or inclined to find a work-around that satisfies the password-strength meter, rather than selecting a truly secure password.

One part of the solution is for users to understand what makes a secure password and apply those standards in their own choices. The judgement of a password-strength meter should not be taken as a final verdict. Rather than placing the burden solely on the user, however, more rigour could be applied in existing meters and in the design of new systems, to support and reinforce choices that actually increase security. Thorough password-strength meters go beyond character complexity and length rules to consider patterns and common passwords from leaked lists, other languages and cracking tools. Those designing password-strength meters don`t need to start from zero; the open source meter implemented by Dropbox is available to be adapted and applied.

Password strength meters are a helpful tool that can reinforce good password choices but the current inconsistency between service providers sends mixed messages.

de Carnavalet, X., & Mannan, M. (2014)"From very weak to very strong: Analyzing password-strength meters" In Network and Distributed System Security Symposium (NDSS 2014).

Foreground Trust as a Security Paradigm: Turning users into strong links

Designing strong tools is an important step in providing security. Unfortunately even highly secure systems can be compromised by human error. Until users understand the 'how' and 'why' of security, they may forever be cast as weak links in the security process. Fortunately human error can be reduced, when systems are designed with the user in mind. When concepts of trust and comfort are leveraged in the design process, a system can support an engaged user to make informed decisions about security. The Ten Commandments for Real People, proposed in an earlier article by the same group of researchers, can guide the development of human-focused security methods and tools (see sidebar).

Marsh et al. posit that raising the profile of the discussion of trust in computing will permit users to make their own security decisions. Foreground trust integrates reasoning techniques into technology, such as in the Device Comfort model. With the Device Comfort approach the technology compiles evidence about the degree of comfort with the environment, the user and their task. The device can then present this information to the user, to advise, encourage and even warn the user about potential actions, in a way helps the user understand the security implications. Device Comfort can recognize context, so permits the user a full range of personas. For example, the parameters can handle the shift from personal use to professional settings, where the user might engage a new set of tasks on different networks. When decisions about security are made in collaboration with the user, rather than unilaterally and behind the scenes by the device, the user is better able to see themselves as part of the security process.

The interface takes the shape of an 'annoying technology'; when an undesirable action is attempted – one that makes the device uncomfortable – the device can (charmingly) get in the way of completing the action. Picture a confirmation step that needs a few more clicks or input from the user. The required action can be more or less involved, depending on the level of device comfort. Rather than block or override the user instruction, the device asks the user to pay attention before completing the action. This momentary obstruction brings to the foreground that the user is making a security decision. The device is enabled to, in turn, empower the user to judge their level of trust.

Ten Commandments for Real People

1. The model is for people.
2. The model should be understandable, not just by mathematics professors, but also by the people who are expected to use and make decisions with or from it.
3. Allow for monitoring and intervention. Humans weigh trust and risk in ways that cannot be fully predicted. A human needs to be able to make the judgment, especially when the model is in doubt.
4. The model should not fail silently, but should prompt for and expect input on 'failure' or uncertainty
5. The model should allow for a deep level of configuration. Trust and security models should not assume what is `best' for the user. Only the user can make that call.
6. The model should allow for querying: a user may want to know more about a system or a context.
7. The model should cater for different time priorities.
8. The model should allow for incompleteness.
9. Beware your context.
10. Acknowledge fragility.

Security concerns are not going away, but neither are users. In designing and commissioning security software, consider the Ten Commandments for Real People. Rather than relegating the user to 'hopeless' status, foreground trust returns some authority to the human. Security is an ongoing process and one in which people can benefit from the reasoning ability of machines to make smart decisions about security.

Humane security decision-making assistance by technology can improve security decisions by humans.

Marsh, S., Dwyer, N., Basu, A., et al. (2014). "Foreground Trust as a Security Paradigm: Turning Users into Strong Links." Information Security in Diverse Computing Environments, 8.

Cutting the Gordian Knot: A Look Under the Hood of Ransomware Attacks

Malware is old news, but recent years have seen an increase in the type known as ransomware. Ransomware locks the affected computer and demands payment from the victim in order to restore their files. Once infected, the victim has few choices but to pay the ransom. In this way, ransomware presents a Gordian knot - a metaphor for a complex problem, one that cannot be 'untied' using conventional thinking, so must be 'cut'.

Kharraz et al. examined characteristics of the most common types of ransomware. The 1,359 samples from 15 'families' of malware cover most of the ransomware observed in the real world between 2006 and 2014. The researchers allowed the malware to run in a controlled laboratory test environment, in which they can limit damage and also monitor the malware behavior. The traits observed in the lab can inform design of detection mechanisms. Fortunately, defending against attacks is not as complicated as previously assumed.

Most ransomware are not very sophisticated. Many types lack the ability to execute a comprehensive attack; they rely on superficial approaches to disabling user access, such as by locking the desktop, rather than more thorough methods like encrypting or deleting files. Several different classes of ransomware cause similar changes in file system activity. Because this pattern is different from the activity of benign – or non-malware – processes, malware activity stands out among file system requests. For example, a system that receives a large number of similar encryption or deletion requests might be infected. A useful protection capability would be to intercept file system requests, then discard or confirm suspicious requests before implementing them. This early detection would also allow an earlier intervention to recover deleted files.

The malware may not be overly complex, but an analysis of ransom payment data suggests cybercriminals who use ransomware have been evolving their financing methods to avoid detection. In order to make tracing more difficult and to improve the portability to other currency, payments are made using Bitcoin, often in small amounts, and receiving Bitcoin addresses are used for only a few transactions over a short time period. The attack software may be simplistic, but the financial side is increasingly deliberate and sophisticated.

There are ways to detect and limit damage from ransomware attacks. With an improved knowledge of the patterns and processes involved in executing a ransomware infection, system managers can focus on improved monitoring and so enhance defenses against this particular malware.

Intercepting suspicious patterns of file system requests would greatly limit the effectiveness of the worst forms of ransomware.

Kharraz, A., Robertson,W., Balzarotti, D., Bilge, L. & Kirda, E. (2015) "Cutting the Gordian Knot: A Look Under the Hood of Ransomware Attacks." In Detection of Intrusions and Malware, and Vulnerability Assessment: DIMVA2015 Proceedings

A Human Capital Model for Mitigating Security Analyst Burnout

Many organizations dedicate significant resources to security monitoring. The ability to identify threats to security depends on effective analysts so it's troubling that this group is often characterized by high rates of burnout. Although there is clear concern about the consequences of frustration and fatigue, such as increased risk of poor judgment and turnover, little is known about what's behind the phenomenon.

Sundaramurthy et al. worked with the security analysts from one Security Operation Center (SOC) to identify the factors behind security analyst burnout. Using an approach long accepted in anthropology, known as Grounded Theory, one researcher was embedded as an analyst in a corporate SOC – both doing and observing the job on a daily basis, with minimum interruption to the operations. The researcher took notes on their observations, which were then coded and grouped into a coherent set of findings. In this case, the findings that emerged point to analyst burnout as a human capital management issue, stemming from the interaction of human, technical, and managerial factors.

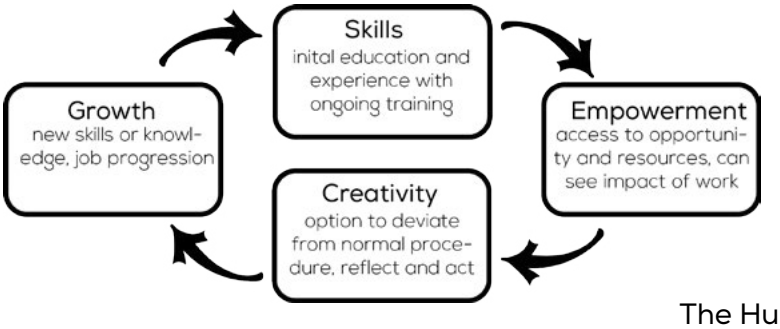

The Human Capital Cycle

Human capital is the sum total of the knowledge, skills, and experiences of the individuals and the team. This asset must be carefully managed, alongside the other resources and operational demands of any security monitoring mandate. The human capital cycle illustrates the factors that reinforce one another in a cyclical manner. Importantly, the reinforcement can be either positive (virtuous) or negative (vicious); when skills are low, for example, an analyst might be only marginally empowered, with few creative outlets and growth opportunities. Someone with a higher level of skill, however, might be granted more opportunities, so experience an increase in their knowledge and skills.

There are several factors that intersect with this model in shaping security analysts' enthusiasm for their work. Done right, these factors can reinforce the virtuous cycle; however, mismanaging any of these factors could contribute to the negative reinforcement of a vicious cycle.

Operational efficiency can be realised through adequate and clear communication to support cooperation. This is particularly important in organizational contexts where the analyst's workflow is dependent on other departments with different priorities. Automation can be empowering when based on true needs identified through reflection, can permit the redirection of the analyst's energy to less repetitive tasks and also provide a creative outlet. Performance metrics can contribute to satisfaction where measures are seen as reflecting analysts' achievements. However, these metrics must not be relied upon for people or process management but integrated with other analyses with the understanding that metrics oversimplify and cannot capture the full extent of operational activities.

Analyst job satisfaction is integral to a sustainable security enterprise. To effectively and proactively address the issue of analyst burnout, those who select or manage security monitoring operations can watch for opportunities to have analysts contribute to planning about operations and automation, as well as decision-making about meaningful metrics. Whether in-house or externally provided, security operations must have access to effective communication and smooth workflows between departments.

Human capital management can reinforce a virtuous cycle of improvement for security analysts and prevent burnout.

Sundaramurthy, S.C., et al. (2015). "A Human Capital Model for Mitigating Security Analyst Burnout." Presented at the Symposium On Usable Privacy and Security (SOUPS).

Robust Fingerprinting for Relocatable Code

A large number of security and forensic applications, especially in a cloud environment, rely on the fingerprinting of executable code. There are many ways of combining lines of programming instructions or code to achieve a result. This creates unique patterns of code, which much like a fingerprint can be used for more reliable identification. The goal of this study is to support security and management services in large-scale virtualized environments by enhancing the ability to identify code regardless of where it is .

Ahmed et al. present primary research on a technique to identify a known piece of code that is running in an arbitrary location. The authors developed a research prototype, and evaluated its effectiveness on more than 50,000 samples of executable programs.

The authors recognized that previous fingerprinting techniques for Microsoft Windows were not able to work with partial information and required an expert analyst to drive the discovery process. To overcome this problem they proposed a new solution using relocation tables (a directory of where each executable code is located) as the key identifying characteristic. Relocation tables have inherent patterns that are quite distinct, which can be used accurately and efficiently to identify known executable code.

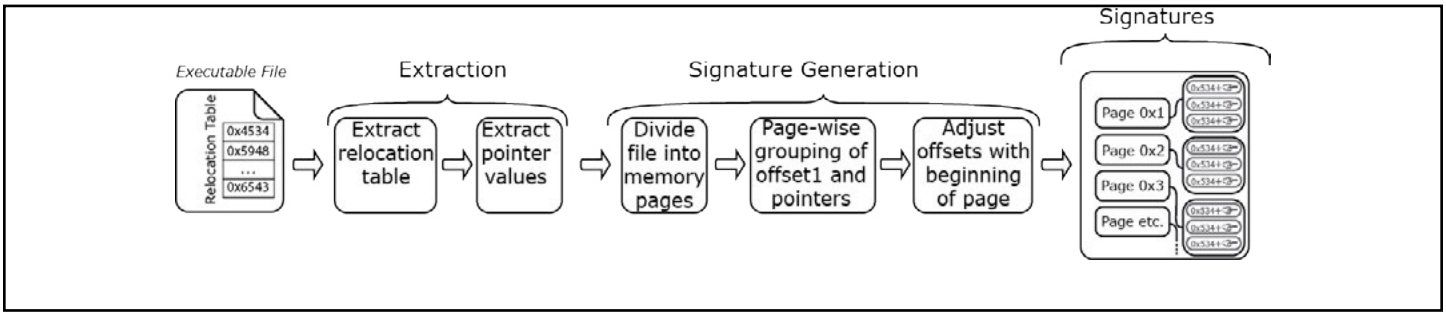

Relocation Table Focused fingerprinting

The method starts by dividing an executable file into smaller blocks, with each block given a generated signature (fingerprint). A block signature consists of relocations in the code and the offset values for the relocations. These signatures are quite unique and can be easily identified in the relocation tables. The only input required to generate a signature is the file containing the executable code. Empirical results show nearly 100% accuracy when identifying fingerprints. Essentially, this technique and tool can find a segment of executable code that has been moved by identifying the fingerprint left in the relocation table.

Identifying executable code is an important aspect of security monitoring and forensic analysis. The research provides another tool for code identification, which will become increasingly important in complex environments and investigations. Although malware detection and classification was not the purpose of this study, the method can be applied to finding known malware executables.

There is a growing trend towards cloud providers offering ever more sophisticated security-related services for their tenants. This research could provide enhanced abilities to providers by offering a robust method for automated fingerprinting.

Fully automated and feasible code identification through fingerprinting is possible, providing promise of new tools for security professionals.

Ahmed, I., Roussev, V., & Ali Gombe, A. (2015) "Robust Fingerprinting for Relocatable Code" in Proceedings of the 5th ACM Conference on Data and Application Security and Privacy 219-229

Reliable Old Wineskins: The Applicability of the Just War Tradition to Military Cyber Operations

The relatively new phenomenon of attacks and warfare in cyberspace make it difficult to understand the effects of these actions. The invention of the 'Cyber Domain' provides very different warfare capabilities and challenges to previous technological advances. While international law has established some rules and regulations for cyber attacks, it may be that traditional moral frameworks still apply in cyberspace.

Barrett's research builds on the existing legal framework for military cyber operations to include an ethicist's perspective on the usefulness of including both law and the traditional just war doctrine to provide ethical guidance. The just war theory provides ethical criteria for decision-making about going to war and taking actions during war.

The Just War Doctrine.

International law leaves a number of ethical and moral questions unaddressed for conducting military operations in cyberspace. For example, should contractors engage in military cyber operations or should autonomous cyber weapons exclude a human decision maker?

Barrett argues that both contractors and autonomous cyber weapons would not be able to act fully in accordance with the just war criteria of discrimination and complete attribution. Because of this contractors should not be deemed proper authority and autonomous weapons should be cautiously pursued. Barrett asserts that international legal definitions can lead to overly permissive and overly restrictive conclusions. The Tallinn Manual, a guide on how international law applies to actions in cyberspace, allows that operations which do not use "force" be directed against civilians and their objects. While international law may permit a cyber attack against civilians, Barrett argues that directing attacks against civilians would be unjust.

The just war criteria stipulates that only a proper authority (most likely a nation-state) can respond to a threat against life and life-sustaining properties in self-defence. A defensive response must be proportional and in defence of a real threat to life. While the proper authority has a right to self-defence, the authority also has to follow state responsibility. This means that if an attack originates from within its boundaries, the state must act to quell the attack (even if the attack is directed at another state or their citizens).

Difficulties with attack attribution in cyberspace create a challenge in identifying whom to respond to. To be deemed a justified defensive response, the attack must be attributed with absolute certainty, not just circumstantial evidence. If involving a state, the state responsibility cannot be assumed. If involving an innocent third-party (not responsible for the attack, but the attack originated in or routed through this state), the third-party state has sovereignty rights, but is morally obligated to protect other states' citizens from the threat.

Furthermore, the defensive response must be able to discriminate the target, which is difficult in this highly unpredictable environment. While cyber weapons may seem more ethical than kinetic ones, users must exercise due diligence - adequate weapons testing and intelligence to ensure the impact is directed specifically where intended.

International law is insufficient in guiding cyber conflict, however where combined with a moral agent operating within the just war criteria are capable of providing the ethical guidance needed to legitimately conduct military cyber operations.

The Just War Doctrine provides ethical guidance on cyber security decision-making that can compliment international law.

Barrett, E. T. (2015). "Reliable Old Wineskins: The Applicability of the Just War Tradition to Military Cyber Operations." Philosophy & Technology, 1-19.

Experimenting at Scale with Google Chrome's SSL Warning

Modern browsers warn users when they are potentially clicking from an encrypted connection, using the Secure Socket Layer (SSL) to an unprotected connection. This can indicate a potential threat to the user. Despite the difficulty in assessing the risk highlighted by the warning, many users 'click-through' this message. Past research has shown that more people using Google Chrome ignore warnings than those using Mozilla Firefox. This raises the question of how the security warning itself affects the user assessment of risk.

Felt et al. set out to understand why Firefox has a 37% lower 'click-through' rate on SSL warnings than Google Chrome. The researchers designed six experimental warnings to test four hypotheses about the impact of visual design, an extra warning stage, corporate style guidelines, and the image of a watching human on lowering click through rates (CTRs). The experimental warnings were included in Google Chrome 29, providing differing warnings to those opting to 'send crash reports and statistics to Google'. The researchers measured 130,754 impressions under field conditions.

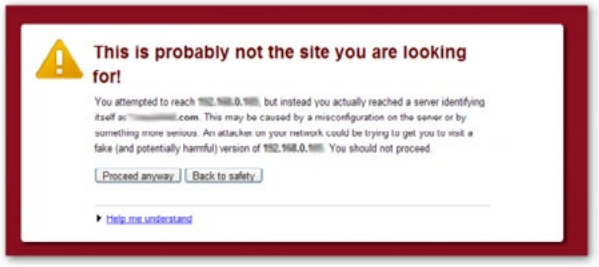

The Google Chrome SSL Warning

The Mozilla Firefox SSL Warning

The study provided some interesting outcomes. In a test where users had to click through a warning twice, the second did little with almost all participants (98%) clicking through both. The use of corporate style or branding had little effect and the use of a watching figure, such as a police officer did not increase the effectiveness of the warning.

The results of the field study suggest that the visual appearance of the SSL warning accounts for one-third to one-half of the difference in effectiveness between Firefox and Google Chrome. Therefore, the design of SSL warnings can drive users towards lower click-through rates and thus safer decisions.
This research contributes to the ability of SSL warning designers to create effective warnings that avoid the use of technical jargon by de-emphasizing technical details, include information on ways to mitigate risk, and include a clear default choice.

Users can be encouraged to make safer decisions by modifying the appearance of security warnings.

Felt, A. P., Reeder, R. W., Almuhimedi, H., & Consolvo, S. (2014) "Experimenting at scale with google chrome's SSL warning". in Proceedings of the 32nd annual ACM conference on Human factors in computing systems. 2667-2670

Owning Your Home Network: Router Security Revisited

Digital Subscriber Line (DSL) routers have become an essential part of home networks as they provide the gateway to broadband Internet. This gateway function also makes these devices important from a security standpoint. As the frontline of increasingly sensitive home networks, DSL routers become the subject ofan important question. Are they secure enough or can DSL manufacturers make changes toimprove the security of all home networks ?

Niemietz and Schwenk conducted a study on the security of home network routers. The researchers chose routers from ten different manufacturers, based on Amazon's top ten most popular list, to study the security vulnerabilities of these inexpensive products.

To analyze the security of the routers, the researchers began the study by launching a number of different attacks at the router's Web interface. The goal of each attack is to gain full control of the router and the user's network. Once full control is gained, the attacker can change critical settings, make the network unavailable for a specific time and use the access to build botnets.

They used the techniques of Cross-Site Scripting, Cross Site Request Forgery and User Interface redressing to manipulate the victim to click on a malicious link or unknowingly share their user and password information.

In the process of analyzing the vulnerabilities of home routers, this study found that:

- All tested routers were vulnerable to Web-based attacks, allowing the researchers to gain full access to the network through the router
- All tested routers had identical default passwords
- Eight out of ten routers were vulnerable to Cross Site Script (XSS) attacks
- None of the tested routers had protection mechanisms in place against User Interface (UI) redressing

By proving that cheap routers can be easily hacked, the researchers were able to produce a list of simple countermeasures for manufacturers to use in enhancing the security of their product.

- Create randomly generated default login data. The attacker has a 55% chance of guessing the right login data due to common default passwords.
- Minimize information leakage by using a randomly generated string for default naming printed on a label (ex. Router Login XXX rather than TP-Link WR841N)
- Use only signed SSL/TSL certificates from a certification authority
- Validate User input on interfaces as many attacks rely on input sinks, which are not properly validated
- Use X-Frame Options to protect against attacks using iFramea
- Set the window.name variable to a random value
- Flag cookies

The implications of this study are not limited to DSL routers as many other devices share the web interface internet connectivity feature and possible vulnerabilities such as: network switches, smart TV systems, network-attached storage devices, etc.

This study gives a representative overview of the security of current home router Web interfaces, and the possibilities to gain full access through Web-based attacks. The flaws of these systems put home networks at risk. The countermeasures listed are well known; therefore this study is a necessary security evaluation for manufacturers and purchasers of DSL routers.

Home DSL Router security is poor and can be greatly improved by small changes made by manufacturers.

Niemietz, M. & Schwenk, J. (2015) "Owning Your Home Network: Router Security Revisited." presented at the 9th Workshop on Web 2.0 Security and Privacy (W2SP), 2015.

SERENE-RISC Six Key Activities

Smart Cybersecurity Network (SERENE-RISC) is a knowledge mobilization network that organises six key activities intended to reach its various audiences: workshops and seminars, a knowledge brokers' forum, quarterly knowledge digests, Konnect - online knowledge-sharing platform, a public website and a professional development program.

The latest developments and upcoming activities for each of the six key activities are shown below.

The SERENE-RISC Quarterly Cybersecurity Knowledge Digest
2015 Summer
Editor-in-Chief Michael Joyce; **Editors**: Emily Maddocks, Justin Anstett; **Scientific Editor** Benoît Dupont

To receive the latest issue and access back issues apply for a free membership at info@serene-risc.ca

Government of Canada / Networks of Centres of Excellence
Gouvernement du Canada / Réseaux de centres d'excellence
serene risc
Université de Montréal