Synthesis Report on SuRI Workshop on Cyber Risk & Information Security held on June 2-4, 2014

IMPROVING CYBER AND INFORMATION SECURITY

Anjali Nursimulu, IRGC

Cyber technology is becoming increasingly pervasive, driven in part by the increase in low-end embedded technologies, for example, in household appliances and devices in order to provide new functionalities. The average number of IP-based devices in personal use was 0.1, 4 and 140 in 2007, 2010 and 2013, respectively. This growth enables higher interconnectivity, supporting the emergence of such trends as Big Data and the internet-of-everything. Together with the increase in malice and asset lifecycle mismatch, the potential for mounting privacy and security risks is becoming evident.

Foreseeable risk-o-meter

Developing a fix to the problem is difficult. Academic research does not always translate into practical strategies or policies. In the absence of liability, Industry does not find business value in addressing the emerging risks. And, intelligence agencies focus on national threats and surveillance. The need for policies and guidelines to address shortcomings in privacy and security strategies in view of such practices as bolted-on security and the rise of malice is unequivocal.

THE CHALLENGE OF ADDRESSING TARGETED VERSUS BENIGN FAILURES[1]

Targeted cyber-attack is the new challenge in IT security. Marketability of information — for targeted advertising, insurance, banking and even financial trading as in Bloomberg sentiment analysis tool — is fuelling the growth of espionage and the concurrent growth of an enabling black market.

Malice is dynamic, adaptive and reactive to changes in practitioner's product and defense strategies. The practitioner therefore needs to be proactive and not reactive. Probabilistic risk assessment, while relevant for random failures, is not sufficient or simply does not work in malicious environment. Traditional probabilistic threat models fail to address the systemic and persistent nature of targeted malice. Trojan, FLAME, DUQU, STUXNET attacks for instance are hard to detect and may require frequent resets. But how often and at what cost? How to do so in complex distributed networks?

The business community has yet to realize how much change to operational environment and in design is needed in face of malice. Current engineering practices are maladaptive and misaligned in that what works well in benign-failure environments is flawed in malice environments. Moreover, security is pushed aside because of cost and lack of skills; so risks are passed on to end users. Insufficient emphasis is put on mutual suspicion and other security primitives. Taken together, badly used security models, poorly understood and little analyzed threats constitute the perfect recipe for disaster.

Driven by quick paths to profits rather than the desire to build a quality system, the fast pace of contemporary technological changes are leading to a complex and vulnerable cyber and cyber-physical architectures. Design focuses on ease of use and added functionality rather than security, potentiating technical debt. Technical debt can be due to conceptual errors in the design of a product; this is best solved early on during design phase. There are also implementation errors in the product as built and discovered in deployed products. Thus, even robust products may entail a transfer of debt from vendor to client and consumers. The risk is that such technical debt can suddenly burst.

Relatedly, there is a cyber-trust bubble which may originate from a mismatch between human and cyber trust, and associated cognitive dissonance. Human trust depends on identity, role, capabilities and intent; it is limited, does not scale readily, is easily revoked and once revoked, it is not easily recovered. Cyber trust on the other hand depends on identity and usually not much else. It is transitive—it spreads easily and widely—and is hard to revoke. Cyber trust (and certificates) is poorly understood; and risks are created when people are led astray.

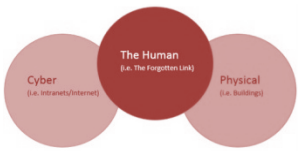

Cyber risks are thus real and need to be studied with further scrutiny. Analysis is required on how fast an event can escalate and with what consequences. Attackers will be always ahead of the game with attacks lined-up. Our society will be vulnerable to cyber failures owing to its complexity. Not all solutions are technical. Threat and security analysis needs to innovate on all three fronts — physical, cyber and human — as they interact.

[1] Presented by Retired NSA'er Brian Snow

PROMOTING SECURITY-BY-DESIGN AND ASSURANCE[2]

The rapid deployment of network-capable appliances in home and office environments raises concerns over the fact that many of these embedded systems have been designed with very simple, if any, threat profile in mind. How trustworthy are they? Conventional antivirus and antimalware will not provide the necessary security. How trustworthy do they need to be? There is social stigma. End-users rarely care if at all they are well-informed about the risks.

It is therefore problematic that emerging consumer-level products typically do not have security designed in. Designing with security in mind as a major requirement only happens in a small number of cases, typically driven by demand for certification or where there are clear liability issues. Low-end embedded devices, such as smart-cards, sensors, RFID, smartphones, peripherals, household devices and many industrial systems, are typically low-cost, low-powered devices with limited facilities and limited memory and may not be capacitated even for secure remote attestation.

Therefore, while scientists are delving into such questions as the smallest number of architectural features that are needed to achieve provably secure remote attestation or the lowest number of additional gates and minimal modifications to connect platforms, forward-looking policy should consider alternative, sustainable solutions.

But, can policy facilitate the large-scale adoption of security-by-design? How should the development and associated assurance costs be funded? How to address the security challenges of software-based legacy devices that cannot be retrofitted? The following section provides some insights from medical device cybersecurity.

BOOSTING CYBERSECURITY OF MEDICAL DEVICES[3]

There is a (flawed) belief that everything is made much safer because of internet. The tightening of medical-device standards by the Food and Drug Administration (FDA)[4] reflects the belief that implantable medical devices should be trustworthy and that improved security will enable medical device innovation. In particular, wireless devices are believed to increase safety by enabling remote monitoring and fine-tuning. But, where does the security bar in the cyber-pervaded and safety-centric medical world really lie? And, what are the real risks?

The main risk is not about hackers breaking into medical devices. The money is not there. The two main risks are wide-scale unavailability of patient care and integrity of medical sensors. There is already evidence that Health Information Technology (HIT) devices can be globally rendered unavoidable. In such cases as in the cases where old software in legacy infrastructures is no longer supported, users are helpless. Manufacturers using wireless, radio, USP port, networking and cloud can no longer claim unawareness of security risks and recognize that software updates should be issued to customers. But how often and at what cost? A related major concern in the field of medicine the mismatch between software and hardware lifecycles, in particular the pervasiveness of legacy and ageing infrastructure, which altogether increase the cost of patching and related risks.

There are also management issues due to diffusion of responsibilities. Software updates involve agency risks such as implementation errors or cumulative effect of unsafe practices such as ignoring update notifications. Furthermore, there is little privacy or security in take-home devices There is also the question of how significant are intentional, malicious malfunctions in software. Factory installed malware is quite common, even if unintentional and accidental. Is there too much software in medical devices? Is there overconfidence in software? What about flaws in design of user-interfaces that induce human error?

From a compliance perspective, it may be relevant to ask engineers what design controls are in place to address cybersecurity risks. But engineers may be protected by limited liability as risks in medical sector are largely due to system level effects and vulnerabilities inherent in the widespread practice of downstream integration of broadly-sourced devices. Security is costly, especially in the absence of liability. Can we rely on trust mechanism for security? If not, there is a need for regulatory policy.

There is a need for premarket evaluation of security and privacy elements of devices and systems and design of post-market systems that enable effective collection of cybersecurity threat indicators for medical devices. But do medical device manufacturers understand threat models whether failures or sabotage? The challenge is that IT people who are responsible for security do not have the requisite safety background and typically compete for airtime. There are new emerging risks with medical devices in the cloud and authentication may not always provide privacy and security protection.

[2] Presented by Andrew Clark
[3] Presented by Kevin Fu
[4] http://www.ihealthbeat.org/articles/2013/6/13/fda-urges-medical-device-makers-to-boost-cybersecurity

IS USABLE SECURITY A UTOPIA?[5]

Auguste Kerkhoff stated almost 130 years ago that the security of a cryptosystem should depend solely on the secrecy of the key, meaning that one should be able to reveal the entire secret of a system to the enemy without causing inconvenience. Importantly the key should be easy to remember and communicate without having to be written down.

But today, there is significant authentication fatigue, driven by such factors as large the number of required resets and complexity of keys. At the workplace, the management of passwords has been documented to increase the workload to the user and typically leads to shortcuts.

These beg the question as to whether it is possible to forget passwords and devise new authentication schemes, focusing on economics — reducing friction, time and workload — instead of user interface. The underlying reasoning is that the problem is deeper than the user interface and instead lies on the incentives and authentication mechanism and the compliance threshold.

As a result, there may be a need for policy, whether internal or externally instituted, to promote security hygiene. This can be facilitated by a security budget to foster prioritization on where authentication is genuinely required and transform existing authentication deployment.

EGRESS ACCOUNTABILITY AS A SOLUTION TO EMERGING PRIVACY CONCERNS?

Massive, systematic loss of privacy is a serious risk. With the growth of algorithmic decision-making, there is scope for vast automated abuse of data that are made available for instance through the digitization of medical records, public posting of genomes, mobile-device tracking and digital finger-printing. In such a hypothetical post-privacy world, it will be impossible to control the information flow.

Privacy-enhancing technologies (PET) such as anonymizers have been poised to provide privacy protection but even TOR anonymity technology has been overridden. Current PET also ignores the background of private data. Egress accountability (EA) is being proposed as an alternative solution. EA is a very general approach — provable without revelation of the algorithm — that ensures that the output of an algorithm is based only on inputs that are considered to be fair.

But in a post-privacy world where privacy is not prevalent, is "certifiable" use of fair-only information a plausible cure? If it is, what regulatory and institutional mechanism will support its deployment? Is it conceivable that the cure is worse than the disease?

CONCLUSIONS: ON THE NEED FOR POLICY

The threat-o-meter is shifting from unsafe practices and accidents to sabotage; cybersecurity risks are now considered foreseeable risks. Security cannot be bolted on; it has to be built in. Design controls in early manufacturing should address risks. Will market lead to more security or do we need regulation?

Owing to limited liability, most manufacturers have little incentive to bear the upfront developmental cost of security-by-design and associated assurance cost or to cooperate on standards, e.g., for interoperability. Depending on firm size, resource constraints may be binding. Likewise, there are lots of variations in CTOs' skills and/or their ability to exercise their capabilities. Left to evolutionary market forces, security outcomes will not change before years even in the presence of net-benefits. Security, as part of the solution, promotes safety, effectiveness, cost reduction, assurance, end-user acceptance, predictability, reliability dependability and, in networks, interoperability.

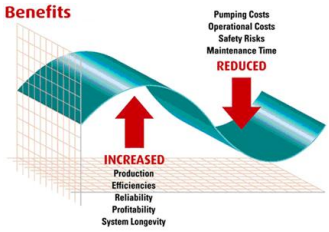

SOME POLICY CONSIDERATIONS

Promote security-by-design:

- Security risk process should parallel safety risk and begin at concept phase.
- Promote and systematize rigorous assurance: installed malware and other implementation errors are not uncommon.
- Promote industry know-how on threat models for complex and engineered cyber risks.

Promote security hygiene and compliance

- Less is more: only cyber elements that have meaningful use should be embedded in our devices; likewise only relevant and fair information should be used. Authentication should be used with parsimony.
- Provide a platform for sharing of threat indicators.
- Pre-market review and cybersecurity updates of COTS software should be systematized.

Incentivize innovative cyber security beyond compliance:

- Probabilistic testing and reputation risk: What is brought to market should fulfill security engineering expectation; the thinking and threat model should be documented in place of check list approaches and manufacturers warned of the forensic scrutiny in case of recall.
- Promote research on and experimentation with novel authentication and privacy-preserving methods.

[5] Presented by Angela Sasse

Report on panel discussion held on June 2, 2014 at EPFL on

CHALLENGES AND OPPORTUNITIES OF AUTONOMOUS CARS

Anjali Nursimulu, IRGC

LOGICAL INCREMENTALISM. Opening the session, Arnaud de la Fortelle, Director of the Robotics Lab at Mines ParisTech, highlighted that the idea of autonomous cars is not new. The concepts of automated road, automated trains and unmanned trams have been around for several decades already. And, automated speed control, lane keeping and longitudinal distance automate control facilities are common-place. In this light, the introduction of autonomous cars is consistent with an evolutionary trend in passenger mobility, and not vulnerable to attacks from the cyberspace as long as such cars are intelligent but not cooperative. The Google car is one such prototype. But the risk of physical cyber attacks or near-field attacks cannot be eliminated. This said, more research needs to be undertaken for fully automated systems.

GIANT EVOLUTIONARY LEAP. John Scott, Chief Risk Officer at Zurich Insurance, remarked that it takes time for cars to change, thus echoing Arnaud's idea of incrementalism. However, he quickly added that autonomous cars provide a tremendous opportunity for designing entirely new mobility concepts in new cities, where the requisite infrastructure for fully networked road-systems, encompassing cooperative vehicles, can be built into the architecture. This is essentially because of the possibility to start from scratch. Masdar City project in Abu Dhabi is one such example. The drastic change requires not only technological advances but also changes in business models from product to service provision and societal change from an ownership economy to a sharing economy. But, in general, the safety — encompassing built-in fail-safe mechanism for risks from carry-in device connectivity—and economics of autonomous cars are believed to drive public acceptance over the tipping point.

STANDARDS AND INTEROPERABILITY. Speaking to the topic of rules of telecommunications for business, Sylvain Glatz, Specialist in Universal Service, Network Neutrality and Quality of Services and Internet of Things at OFCOM/BAKOM, mentioned five points that will be need to be addressed from a regulatory perspective. First, do hardwares embedded in autonomous cars conform to radio-frequency emission standards? Second, should cars have IP addresses? Third, what should be the obligations of new service providers? Fourth, how to design adapted emergency call assistance for autonomous vehicles? Fifth, what will be the golden rules for standards to ensure interoperability without sacrificing security and who will be responsible for writing them?

BUSINESS OPPORTUNITIES FOR NEW PLAYERS.

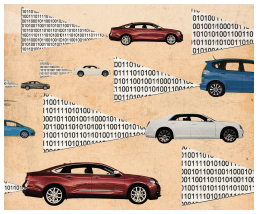

Victor Schlegel, Head of Business Intelligence and BigData Services at Swisscom AG, discussed the opportunities associated with automated cars and automated road traffic for placing new services related to big data on Swiss Market. He substantiated Swisscom's strategy of tapping into new growth areas by presenting models Swisscom has developed to predict traffic flows on motorways based on anonymized mobile use data.

The ensuing discussion was lively and varied. Jean-Pierre Hubaux remarked that with the exception of Spotify and Skype, most **IT innovation comes from US. Will Europe be a laggard in the age of Internet-of-Things?** In reality, Volvo has a fleet of 100 driverless cars — larger than Google Cars. Mercedes, Audi, Volkswagen and BMW have most of the components of driverless cars. One could speculate that these large players are more cautious in revealing their strategies and development compared to Google, for which the Google Car is a mere positioning strategy. But, as with all new technologies there will be winners and losers. Traditional selling points such as " driving experience" will be displaced. The value chain of passenger will expand as new players such as Swisscom, Google and IBM position themselves to provide mobility services for intelligent vehicles.

INTERMITTENT COEXISTENCE AND TRAFFIC RULES.

Speaking on behalf of the industry partners, collaborating with the EPFL Transportation Center, Michaël Thémans discussed two research streams, namely the design of affordable autonomous cars for mass deployment and the investigation of traffic rules for cooperative road traffic in view of planned automation of traffic network for 2030. This research can also inform rules to facilitate the unavoidable and intermittent co-existence of driven and driverless cars. The research also involves in foresight about future traffic flows on highways, roundabouts and junctions. In mixed traffic, are driverless and driven cars friend or foe? How will drivers react to a driverless counterpart on the street?

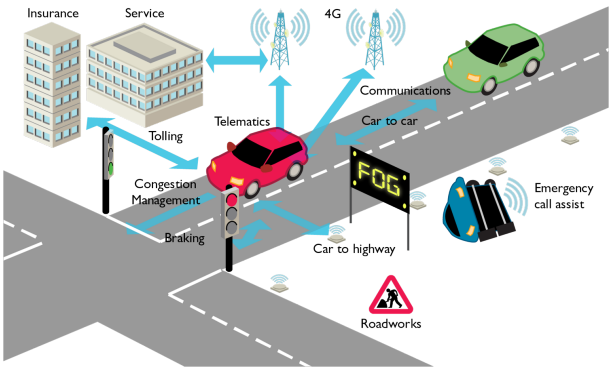

RISK AND COMPLEXITY.

Prof. James Larus, Dean at EPFL's School of Computer and Communications Science raised concern over the apparent trend towards cooperative and intelligent transportation systems. These will involve an increasing number of connections and (near) real-time feedbacks among vehicles and critical infrastructures. The resultant increase in complexity imply a potential for escalation of risks, whether from random failures that cascade through the network or from cyber sabotage. The fully automated network architecture will not only be harder to secure but also costlier.

INSURANCE AND LIABILITY.

Autonomous cars are purported to be safer, reducing the number of accidents and casualties. The risk is reduced—which translates into savings on insurance policy, but who bears the liability in the event of an accident? The current Vienna convention stipulates that the driver has responsility of the car and therefore bears the liability. But does such responsibility hold in the Google Car where there is no possibility of manual override of the automatic system. Who is liable? Is it the car manufacturer? Is it the software designer? What if the autonomous car is endowed with a learning algorithm that results in emergent behavior? Are new liability laws needed? Is mutual fund insurance the way forward?

SOME PENDING QUESTIONS

1. What will be the residual risks (hazard, vulneratbilities, exposures and threats) in a world with fully cooperative road traffic?
2. What mechanisms will need to be put in place to mitigate the impact of such risks in the event of risk-materialisation?
3. What business and insurance models will support the foreseen cooperative traffic?