1. The utopia of security

In any trade, the peddler of goods or services needs to start with a sales pitch promoting their craft. Books on security are no exception, and often start with anecdotal references to security breach incidents. In conformity with that tradition, we commence with a very brief glimpse of the ground realities of the (cyber-)world we live in.

In a 2011 report [1] on Organised Crime Threat Assessment (OCTA), the Europol noted that organised crime groups derived more than 1.5 billion euros from payment card fraud in the EU in 2009. Another 2011 Europol press release [2] reported on Operation Night Clone which led to the arrest of key players and dismantling of one cross-continental racket involved in skimming of EU citizens' ATM/credit cards, which involved law-enforcement agencies across Europe and USA. The particular organization in question was estimated to have caused a damage of 50 million Euros. The article also provided interesting insights on the vulnerabilities that were exploited by the criminals. Introduction of EMV (named after the companies Europay, MasterCard, Visa which created the standard) [3] system (Chip and PIN) across Europe had reduced the risk of skimming of magnetic stipe when the cards were being used within Europe, however, transactions outside Europe, particularly in US and Africa, where EMV was not fully deployed, allowed ample opportunities for the criminals. The above OCTA report estimated that 80 percent of fraud incidents hitting cardholders in the European Union were in fact committed in the U.S.

There are multiple lessons that can be drawn from this single anecdote. A vast and increasing amount of our financial transactions are electronic (cashless) in nature. By some estimates, less than 10% of world's money is currently in the form of physical cash, while the rest is all electronic. Given the ease and portability of use, the trend towards a cashless society [4] continues. Criminals have thus huge interest in exploiting one and all possible vulnerabilities to pilfer, and the damages can be exorbitant, even though the heists do not have the same visual and violent spectacle as the good old bank robberies pictured in classic Western movies of yore.

[1] Europol. EU Organized Crime Threat Assessment. OCTA 2011, 2011a

[2] Europol, 2011b

[3] `http://en.wikipedia.org/wiki/EMV`

[4] Wolman, 2012

In fact, on many occasions the victims may not even realize that they have been robbed, long after the crime has been committed. The geographic spread of the criminal networks make it a law-enforcement nightmare to apprehend the criminals, and even while some groups are brought to justice every once in a while, many more continue to evade the law, and continue with their criminal activities. More disturbingly, even when vulnerabilities are known, and technological solutions may exist to prevent the exploitation of the vulnerabilities, either negligence, or other logistic inertia may prevent deployment of the solutions, and leave the system exposed for further attacks. Thus, technology in itself is not adequate, but the way the technology is used, and the people involved in the process, also play a critical role in achieving or compromising security.

Numerous attacks thus often include social engineering [5] as part of the whole attack process. Ultimately, the adversary would seem to have an upper hand in the game of security, and requires but just one way to exploit some vulnerability, while the security practitioner needs to ensure that all avenues of exposure are well guarded.

This last lesson is encapsulated in the Principle of Easiest Penetration [6] which states that: "An intruder must be expected to use any available means of penetration. The penetration may not necessarily be by the most obvious means, nor is it necessarily the one against which the most solid defense has been installed. And it certainly does not have to be the way we want the attacker to behave."

Another example of the principle of easiest penetration is witnessed from the exploits of the hacker Albert Gonzalez [7] whose modus operandi has been to use SQL injection attacks to deploy backdoors on corporate systems, in order to then launch packet sniffing on Address Resolution Protocol (ARP Spoofing), effectively realizing a man-in-the-middle attack, allowing him to steal computer data from internal corporate networks. There are again several lessons from this incident. SQL injection attack is well-understood, and in theory, easily preventable by proper implementation, yet the vulnerability persists in many old as well as newly implemented systems. The 'flaw' of the ARP protocol which is exploited is in fact a feature used for other purposes,[8] and in principle, an outsider cannot leverage on it. However, because of the successful SQL injection attack as the first stage of the overall attack, in this specific case, the attacker gained privileges of an insider and could thus manipulate ARP subsequently. This shows how there can be unintended consequences of system design choices, and highlights why it is impossible to thoroughly analyze a system's vulnerabilities, or achieve absolute security. And thus, for the adversary to carry out an attack, it becomes a matter of finding one possible way to do so, exploiting

[5] Hadnagy, 2011

[6] Pfleeger and Pfleeger, 2007

SQL injection is a code injection technique, used to attack data-driven applications, in which malicious SQL statements are inserted into an entry field for execution (e.g. to dump the database contents to the attacker). SQL injection must exploit a security vulnerability in an application's software, for example, when user input is either incorrectly filtered for string literal escape characters embedded in SQL statements or user input is not strongly typed and unexpectedly executed. SQL injection is mostly known as an attack vector for websites but can be used to attack any type of SQL database.

[7] http://en.wikipedia.org/wiki/Albert_Gonzalez

[8] The techniques that are used in ARP spoofing can also be used to implement redundancy of network services. For example, some software allows a backup server to issue a gratuitous ARP request in order to take over for a defective server and transparently offer redundancy. ARP spoofing is often used by developers to debug IP traffic between two hosts when a switch is in use.

unforeseen or neglected bugs, or harnessing features in a manner that the original designers did not anticipate.

The Stuxnet computer worm incident [9] is a well known example of exploiting computer and network vulnerabilities to attack industrial Programmable Logic Controllers. In late June 2014, it was revealed by Helsinki, Finland based computer security company F-Secure [10] that a variant of the Havex malware affecting sensitive industrial control systems was being deployed by poisoning legitimate apps. Specifically, the attackers abuse vulnerabilities in the software used to run the websites of companies providing legitimate software to the industrial customers, to break in and replace the legitimate software installers with remote access trojans (RAT).

A remote access Trojan (RAT) is a malware program that includes a back door for administrative control over the target computer. RATs are usually downloaded invisibly with a user-requested program. Once the host system is compromised, the intruder may use it to distribute RATs to other vulnerable computers and establish a botnet. Source: http://searchsecurity.techtarget.com/definition/RAT-remote-access-Trojan

The approach to inject Stuxnet exposed new path, in addition to the established ones like spam/phishing emails, to spread a malware. This is yet another example demonstrating that the attack trajectory could come through unexpected avenues. In this case, vulnerability of the website of trusted vendors was used to gain access to an apparently well secured infrastructure.

The ultimate objective is to damage mission critical and industrial activities, either to wreak financial damage and disrupt economic activities. The Stuxnet attack, for instance, was apparently targeted at disrupting Iran's nuclear program, and hence the anticipated perpetuators are possibly state actors [11].

The US indictment of five Chinese military hackers in May 2014 highlights another potential motive [12] of an attack, namely to steal business critical information and intellectual property, popularly known as cyber-espionage. Organizations, as well as nations, need to be wary against such weapons of cyber warfare, which may be waged by both state as well as non-state or stateless entities.

At this juncture, it is worth pausing to introspect that there is no obvious good or bad guys when we discuss cyber security. Either attacker or the victim may be the bad (or, more pragmatically, the worse/better guy) among the two entities. While studying the subject of security, we thus need to view the topic of cybersecurity with academic detachment, be it because we aspire to defend a computing system, or carry out a successful attack against a well defended system.

The July 2015 hacking of the Italian company Hacking Team [13] is an interesting anecdote demonstrating the duality or ambiguity between the good and the bad, when it comes to computer security. Hacking Team's surveillance remote control system (named Galileo) itself is a RAT malware from the perspective of its victims, even though this is typically authorized by the governments (and agencies) which happen to be Hacking Team's clients. These clients themselves

[9] http://en.wikipedia.org/wiki/Stuxnet

[10] http://www.f-secure.com/weblog/archives/00002718.html

[11] https://en.wikipedia.org/wiki/Operation_Olympic_Games

[12] Note that an indictment only means that the charges are leveled formally, but these are yet to be proven as of the time of writing this article on 1st June 2014 http://www.justice.gov/opa/pr/2014/May/14-ag-528.html.

[13] http://www.wsj.com/articles/hacking-software-maker-gets-hacked-1436223757

sometime have dubious to outright outrageous track record in terms of freedom and protection of human rights. Ironically, the same privacy that Hacking Team product denies its victims is what was breached in the hacking incident which laid bare the source code as well Hacking Team's communication and financial information with its customers and business partners. Early analysis of the source code of the remote control system suggest a further irony and poetic justice - as it appears that the software possibly comes with a kill switch that can disrupt normal functioning of the infrastructure of the clients themselves.

But it is not only nations and enterprises that need to be wary about security in the cyberspace. As an individual, besides the risks of being a victim of surveillance (say, through a malware like the one used by Hacking Team as mentioned above) or financial crime, one could also suffer in other ways, including stalking, harassment, etc. A lot of personal information is explicitly as well as implicitly stored and/or transmitted electronically, and the many digital avenues that enrich our lives, also provide ways for adversaries to encroach upon. The widespread phone hacking carried out by the 'News of the World' [14] is one of the more well documented examples of such a loss of individual privacy, though numerous lower profile incidents occur at an alarmingly high frequency. The July 2015 breach of customer information of online cheating site AshleyMadison.com[15] being one of the latest and more talked about one, given the very sensitive nature of the information.

The emergence and embrace of the Internet of Things by governments, industry players and individuals alike, our life is becoming so much more intermingled with the digital world. With the great opportunities to better our lives, there will also be a natural amplification of the kinds of risks and the effects these will have on our lives. Be it hacking pacemakers [16] to cause fatal damage or exploit medical devices [17] to gain access to health information systems, or gain hostile control over automobiles [18], these are all things that are not paranoid fantasies of a possible dystopia in a distant future, but realities of events that have already come to pass - thankfully, mainly as proof of concepts so far.

These anecdotes make it amply clear the many trappings of security breaches, and why it is important to understand and implement security rigorously, and how the human elements play a very crucial role, on top of the technical aspects. The treatment of the current material is driven by this insight of the augmenting role of human factors in cyber-security, and complements other materials looking at the technical aspects of designing and implementing cryptographic primitives and security protocols, or carrying out (digital) forensic

[14] http://www.bbc.com/news/uk-24894403

[15] http://krebsonsecurity.com/2015/07/online-cheating-site-ashleymadison-hacked/

The Internet of Things (IoT, sometimes Internet of Everything) is the network of physical objects or 'things' embedded with electronics, software, sensors, and connectivity to enable objects to exchange data with the manufacturer, operator and/or other connected devices. The Internet of Things allows objects to be sensed and controlled remotely across existing network infrastructure, creating opportunities for more direct integration between the physical world and computer-based systems. Source: https://en.wikipedia.org/wiki/Internet_of_Things

[16] Boyle, 2012

[17] Grau, 2015

[18] http://bits.blogs.nytimes.com/2015/07/21/security-researchers-find-a-way-to-hack-cars

Trivia: Former US vice-president Dick Cheney's doctor had the wireless connection of his heart's defibrillator disabled in 2007, as a preventive measure against any foul play.

investigations to understand security breaches and take remedial measures.

Terminologies & taxonomy

Often, there is a lack of consensus about some key terms one frequently encounters when discussing security. This arises both from the fact that in many context, there is no practical implication of the subtle differences, and also because if a strict dictionary definition is taken literally, then it will be hard to capture lucidly a wide genus. Thus, despite the lack of an unanimity on the precise definitions, we provide a rough guideline on how to broadly interpret and differentiate some of the common terms.

Data versus Information: To begin with, it it worth noting that data and information do not refer to the same thing, though in many context they may be used interchangably. A fragment of raw data in itself may or not reveal any meaningful information. For instance, the number 42 could be The Answer to the Ultimate Question of Life, the Universe, and Everything, or the number of illustrations used in Lewis Carol's Alice's Adventures in Wonderland, or even the former country code of erstwhile Czechoslovakia, or something altogether different.

For instance, it was reported around the 2013 Christmas that in a data breach, more than 40 million user credit/ATM card PINs were stolen from the US retailer giant Target [19]. (For more details on the data breach incident, see the **Case study: Target targeted through HVAC** at the end of this chapter.) To assuage fears of their customers, the company asserted that the data was encrypted with 3DES, and claimed that the strong encryption prevented the hackers from gaining access to the actual PINs, which can be considered as the actual information in this instance. Furthermore, it was argued that the breached server did not carry the encryption/decryption key, and the company asserted that despite possessing the encrypted data, the hackers will not be able to get hold of the content within. Target's assertions notwithstanding, there was wide criticism for its usage of 3DES instead of AES. Subsequently, many reports of Target customers experiencing strange transactions on their cards have also come to fore, indicating that the attackers did get access to the actual content.

Even though there is no known record of 3DES being really broken, there has been several partial attacks, and hence a different encryption standard, namely AES, is currently considered the state-of-the-art. This is yet again an example where, despite the availability of better technological protection, the actual deployed system's vulnerability is amplified because of the utilization of obsolete technology.

Information & Cyber security: Information security can be broadly described as defending information from unauthorized access, use, disclosure, disruption, modification, perusal, inspection,

[19] `http://www.cnet.com/news/target-encrypted-pins-stolen-but-not-encryption-ke`

recording or destruction (definition derived from Wikipedia [20]). This in turn may naturally include protecting the raw data itself, from which the information is derived.

In order to realize information security, an information system needs to be secured at multiple levels - the hardware itself, the software as well any communication involving the system elements[21]. The security of an information system also depends heavily on people, processes and organizational level factors. Procedures or policies have to be implemented to train people (administrators, users and operators) how to use products, or determine network or even physical access to the system elements.

Note that information security, in principle, includes both digital world as well as physical world. For instance, if someone pieces together shredded pieces of a printout from a garbage bin to get access to a specific information, say a trade secret or tender details of a company, the consequence is as bad as if someone could access an electronic copy of the same document. From this (information security) perspective, one may view the typical treatment of cybersecurity as a subset of the general problem of information security. There is also an increased realisation and emphasis on fusing cybersecurity with physical security, to manage security more holistically. We ought to keep this in mind throughout the discourse, even though most of the following treatment will deal with the aspects of security closely emanating from and influencing the digital/cyber-space.

However, modern cyber-infrastructure serves purposes beyond information centric applications. Organizational workflows, automated business processes, cloud and web services, eGovernance, etc. while heavily reliant on information (and backbone information systems), go well beyond just information. In this broader perspective, cybersecurity subsumes information security. This duality is probably why it is hard to find succinct definition or distinction between the terms information and cybersecurity, and there are debates galore on whether they are the same, or if one is a subset of the other, and so on. Unless pressed with compelling reasons, for the rest of this discourse, we shall use these two terms interchangeably, but we will mainly deal with the aspects of the digital realm, but including security beyond that of information. We will explicitly emphasize on the somewhat divergent interpretations only if and when the situation calls for it.

Computer system security: Computer system security typically refers to security issues that are directly tied to the operating system security. However, over time, the notion of a computing system is blurring in multiple manner. Mobile devices and individual 'things' with miniaturized computers and systems of chips at one end of spec-

[20] https://en.wikipedia.org/wiki/Information_security

[21] By system elements, we refer to any sub-system or components used in the information system. This may include a computing unit, a storage disk, a network switch, etc.

trum, to massively parallelized and distributed computing systems at the other end of the spectrum (e.g., cloud systems, data centers, internet of things) may all be viewed as computing systems at different granularities, running individual operating systems, as well as software to orchestrate the functioning of the overall distributed systems. Security concerns emanating from the operating system perspective, but concerning again any of hardware, software or communication would fall under the broader umbrella of computer system security.

Network versus Internet security: Network security refers to the security of network protocols such as security of TCP/IP, IPSec, OpenSSL, etc. In contrast, the term Internet security also encompasses security issues pertaining and emanating from internet applications - these include browsers, webservers, botnets (portmanteau of robot and network), denial of service attacks, and so on.

A famous example of a bug leading to breach of Network security is the Heartbleed bug in the OpenSSL cryptography library, which is a widely used implementation of the Transport Layer Security (TLS) protocol. The bug (along with a fix for OpenSSL) was disclosed to the public on April 7, 2014. At the time of disclosure, some 17% (around half a million) of the Internet's secure web servers certified by trusted authorities were believed to be vulnerable to the attack, allowing theft of the servers' private keys and users' session cookies and passwords. At the time of disclosure of this 'zero-day' vulnerability, Heartbleed was deemed as one of the most dangerous bugs (in terms of the number of affected parties) in the history.

Application security: Security of and arising from the overlying applications, including from vulnerabilities of the underlying system exposed through flaws of the application fall under the realm of application security. It encompasses measures to be taken throughout the application's life-cycle, including adaptation/patching of the application to prevent exploitation of vulnerabilities inherent to the application, or where the application is a conduit to exploit vulnerabilities of the underlying system.

For example of breach of security at application level, consider the famous Whatsapp messenger, which had deployed encryption in a flawed manner on multiple occasions. This include their initial attempt at encryption using IMEIs (International Mobile Station Equipment Identity) and MAC (media access control) addresses as encryption keys [22] though these information are not secret and relatively easily obtained, then subsequently using one-time pad twice [23].

The Stagefright vulnerability in Android systems revealed in 2015 [24] makes another interesting albeit scary example. Some messaging applications, upon receipt of a multi-media message, but even without the end-user choosing to open the content, would 'pre-process' the media content for faster response time when the user chooses to play the content. However, what would be an optimization trick unfortunately can be exploited by an attacker to execute malicious code, even without any explicit action by the user (in contrast to the typical modus operandi in spear-phishing[25] where explicit action on the part of the user is required, e.g., opening an attachment, etc.). Thus, just knowing the contact information of an intended victim with Stagefright vulnerability is adequate for a successful attack. It

[22] https://nakedsecurity.sophos.com/2013/01/29/whatsapps-privacy-investigated-by-joint-canadian-d

[23] https://blog.thijsalkema.de/blog/2013/10/08/piercing-through-whatsapp-s-encryption/

[24] https://blog.zimperium.com/experts-found-a-unicorn-in-the-heart-of-android/

[25] Spear phishing is the technique to send a message which tricks (possibly by impersonation, or by using another compromised device) the recipient to open a malicious payload or url, leading to execution of malicious code on the end user's computing device and leading to unintended and unauthorized activities - e.g., installing a keylogger to steal personal information, or run a ransomware encrypting user's data and making it unavailable, or send the malicious payload to the contact lists in the compromised device, etc.

was estimated by the discoverers of the vulnerability that at the time of its disclosure, the vulnerability affected 95% Android devices, which would translate to roughly 950 million devices.

A system integrity or security breach may be caused intentionally (man made) or may happen accidentally. Mechanisms to prevent such situations, but also cope with them if they come to pass are required, and all fall within the general umbrella of security management. However, a bulk of our discussion will revolve around man-made, deliberate attempts to breach security, though the other scenarios too will be discussed as and where appropriate.

The rest of this section accordingly focusses on intentional attacks leading to security incidents. The discussion is well summarized in Figure 1 derived from [26]. Though somewhat dated with respect to the evolution of technology, the figure still successfully captures the essential taxonomy of security incidents.

[26] Howard and Longstaff, 1998

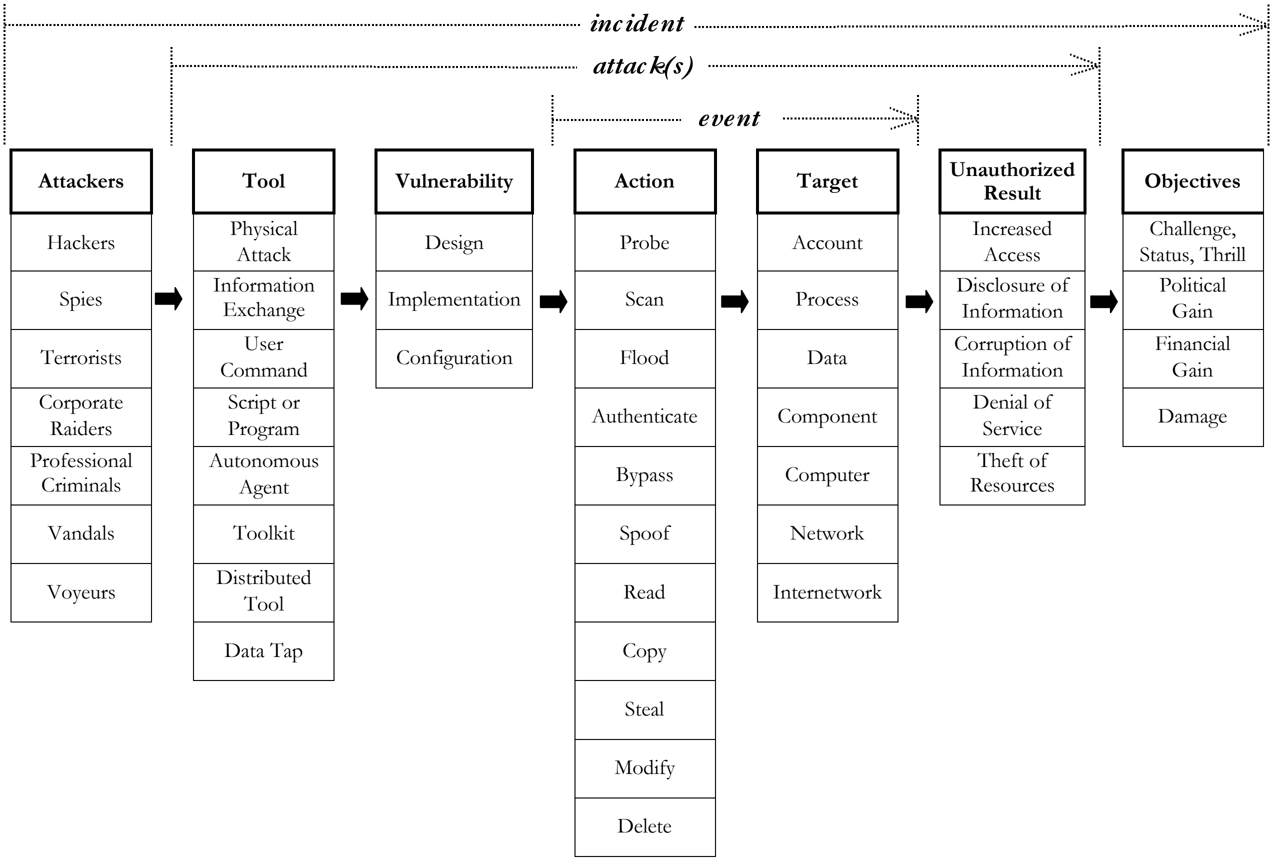

Attackers	Tool	Vulnerability	Action	Target	Unauthorized Result	Objectives
Hackers	Physical Attack	Design	Probe	Account	Increased Access	Challenge, Status, Thrill
Spies	Information Exchange	Implementation	Scan	Process	Disclosure of Information	Political Gain
Terrorists	User Command	Configuration	Flood	Data	Corruption of Information	Financial Gain
Corporate Raiders	Script or Program		Authenticate	Component	Denial of Service	Damage
Professional Criminals	Autonomous Agent		Bypass	Computer	Theft of Resources	
Vandals	Toolkit		Spoof	Network		
Voyeurs	Distributed Tool		Read	Internetwork		
	Data Tap		Copy			
			Steal			
			Modify			
			Delete			

Figure 1: A taxonomy for security incidents.

Attackers could have varied motivations and objectives while carrying out an attack. It could be mere thrill seeking on the part of the attacker, or the intent may be to make a political statement, sabotage an organization's activities, cause terror, or make financial gains, to name a few prominent ones. The attacker accordingly could be an individual hacker, or professional organized criminal syndicate, terrorists, spies, state-sponsored organization, etc. The kind of attacker in turn determines the amount of resources they may have at their disposal, as well as the target and intensity of an attack - random, sporadic crime of opportunity or directed and determined.

Different tools can be used by an attacker, including standard, ready made tools targeting typical and known vulnerabilities, to sophisticated ones made by the attacker for specific purpose or target. Vulnerabilities in a targeted system can be classified essentially in three categories. The design of the system could be flawed. The use of IMEIs (International Mobile Station Equipment Identity) and MAC (media access control) addresses as encryption keys in Whatsapp was such a design flaw in the protocol. The Stagefright vulnerability in Android is another such design flaw. The ARP spoofing carried out by Albert Gonzalez also exploited a vulnerability intrinsic to the design (feature) of ARP.

Even if the system design is not flawed, the implementation may be flawed. The Heartbleed bug, resulting from improper input validation is an example of such flawed implementation, which rendered OpenSSL implementation of transport layer security (TLS) vulnerable. SQL injection attacks also fall in the same category (flawed implementation). Use of an easy to guess password in an otherwise well designed and implemented system is a simple example of poor configuration leading to vulnerability in a system. Another example of poor configuration would be to set up access control rules which are not well thought out, allowing access of information to people who do not need/ought to have access to a piece of information.

The actual act of exploiting a vulnerability using one (or multiple) tool involves a set of actions - this could be simply eavesdropping on communication as would have sufficed with the Whatsapp communication when one time key was being used for communication on both directions, or could be an active attack like spoofing (man-in-the-middle) attack, as was carried out on ARP protocol by Albert Gonzalez. Many other kind of actions are also possible, some of which are enumerated in Figure 1. These actions are on specific resources in the system, be it data or a network element, etc., culminating in an unauthorized result. For instance, in the incidents involving Albert Gonzalez that we discussed previously, data from internal corporate network was exposed. From the same incident, we can see that multiple vulnerabilities were exploited using different means, resulting in several events (SQL injection attack, ARP spoofing) each of which resulted in some unauthorized results, which all put together comprised of the overall security incident. In this incident, the ultimate objective for the attacker being financial gain.

Security objectives

Our discussions so far have emphasized the need for security, and used security incident anecdotes to elaborate several aspects of said security. However, the term security itself is not specific enough to elaborate adequately what specific protection is desired. The most intuitive and layman expectation in terms of security may be privacy centric, regarding the confidentiality of information. However, while a denial of service attack on a server does not divulge any secret information to an unauthorized entity, and hence the confidentiality of the information is not violated, it nevertheless falls within the purview of security. There is thus a need to more explicitly enumerate different security goals or objectives, alternatively also called security properties or security attributes. We next expound a few important models that help discern distinct security objectives.

CIA TRIAD: The CIA (confidentiality, integrity and availability) triad has been a cornerstone in encapsulating information security needs.

Confidentiality: Confidentiality is the property, that information is not made available or disclosed to unauthorized individuals, entities, or processes (quoted from ISO 7498-2:1989[27]). Different means to achieve confidentiality could include encryption, access control through login/password, etc.

Integrity: Data integrity refers to ensuring correctness and completeness of data throughout its lifecycle. Integrity may get violated either accidentally - e.g., corruption of a storage disk, or because of lack of a proper implementation (e.g., transactional guarantees such as ACID (Atomicity, Consistency, Isolation, Durability) property being not properly realized), or even intentionally, for example, someone deliberately modifying the data in any undesirable manner, may be, to say, carry out accounting fraud, or defacement of websites, etc.

Availability: The data (or service) needs to be accessible when it is needed. That is the crux of the availability property. Availability may be compromised due to various reasons such as system component failures (be it hardware or software), overload, or by the system being subjected to a denial of service attack to name a few.

Though originally proposed in the context of information security, and the consequent objectives are somewhat data-centric, the CIA objectives naturally extend to all aspects of cybersecurity. As cardinal as confidentiality, integrity and availability attributes are, they also however, by themselves pretty inadequate in holistically identifying desirable security goals. An exhaustive list is impossible to draw, but there are several other popular models and principles, which we study further next.

Starting on 27th April 2007, Estonia, during an ongoing political disagreement with Russia, experienced denial of service attacks (as well as a few incidents of vandalism) at an unprecedented scale, affecting the functioning of major government, communication and broadcasting and financial institutions. This is an example of an attack compromising mainly the availability of the systems and services. As recently as March-April 2015, China has been accused of carrying out massively distributed denial of service attacks (Great Cannon) on websites that made available content of websites which are censored in China (by the other censorship infrastructure called 'Great Firewall').

[27] ISO7498-2:1989, 1989

McCumber's Cube: In 1991, John McCumber proposed an extension to look beyond just security attributes, and look at other dimensions of security, including the state of the information as well as scrutinize the safeguards in place to realize the security attributes under these different states of the information, to provide assurances on the achieved security.

The original proposal looked at the same CIA triad as the desirable security attributes. However, it explicitly distinguished the states of the information being: (i) data at rest (storage), (ii) data in transit across information (sub)-systems and (iii) data under processing. For instance, while cryptanalysis of a basic implementation of the RSA (with large enough a key size) remains infeasible, side-channel attacks [28] are relatively easier (though mitigations against known side-channel attacks exist).

The other pertinent insight, motivating the final dimension in McCumber's Cube is that, security needs to be seen both from a functional perspective (what should the system do?) and an assurance perspective (how do we ensure that the functional requirements are indeed achieved?). The assurance comes through a dimension of safeguards. A quick realization when viewing security from the assurance perspective, is that technology, while essential, is in itself inadequate in order to achieve security. A multi-pronged approach is needed instead. The other crucial safeguards articulated by McCumber include policy & practices and human factor.

The policies & practices aspect looks at determining best practices that should be followed, or standards that need to be adhered to, among others. For instance, it may set out the roles of individuals in an organization, determining who should have access to which set of information, or determine a course of action in the event of a particular incident, and so on.

It is one thing to have a set of policies and practices, but it is a totally different ball game to ensure that the personnel in fact are aware of and adhere to the same. The human factor essentially deals with this, as well as educating them against any other common pitfalls - for instance, not to share with any generous but orphaned Nigerian prince or deposed dictator's widow one's personal banking details.

Parkerian Hexad: In 1998, Donn B. Parker added three further security attributes to the original CIA triad. The six together are known known as the Parkerian hexad. The additional attributes being possession/control, authenticity and utility. These six attributes are deemed atomic, in that they cannot be broken down into other sub-attributes, and they are also non-intersecting, defining unique aspects

[28] In cryptography, a side-channel attack is any attack based on information gained from the physical implementation of a cryptosystem, rather than brute force or theoretical weaknesses in the algorithms (cryptanalysis). For example, timing information, power consumption, electromagnetic leaks or even sound can provide an extra source of information, which can be exploited to break the system. Some side-channel attacks require technical knowledge of the internal operation of the system on which the cryptography is implemented, although others such as differential power analysis are effective as black-box attacks. (Quoted from https://en.wikipedia.org/wiki/Side-channel_attack.)

of security. Though they refer to unique aspects, it is worth noting that they are not necessarily orthogonal, and may have bearing on the other attributes.

Possession: The ownership or control over a specific information or resource. For instance, if encrypted data is stolen then the actual content and its confidentiality is not immediately violated. However, the data owner cannot influence further actions on that data by the attackers in any manner, and depending on the quality of encryption, or the attacker's ability to guess the password, the content may still be exposed at a future time point.

Authenticity: Authenticity refers to the veracity of (authorship of) a piece of information. A successful man-in-the-middle spoofing attack is a classic example where the authenticity of a received communication is violated. Another typical example where authenticity is violated, is any confidence trick based scam such as spear phishing. Public key cryptography based digital signatures is one means to meet authenticity requirements in some setups.

Utility: Utility refers to the usefulness of the information (or resource). For instance, if data is stored encrypted, but the en/decryption password is forgotten, then even if the encrypted data stays available and no violation of confidentiality, integrity, possession or authenticity occurs, the owner will not be able to use the data anymore. Note that in this example, the password information is 'unavailable', so is the 'unencrypted' content. This is a scenario where, as mentioned above, the hexad attributes, while referring to unique aspects, are not necessarily orthogonal. However, they indeed remain unique aspects, because utility can be breached in other manner. For instance, a mere misuse of measuring units/conventions at NASA resulted in the loss of a US$125 million Mars orbiter[29] – an anecdote signifying how the utility of the information was compromised, but availability or other attributes were not.

Absolute security is an unachievable utopia, and consequently, no such list can be exhaustive. One may continue to add further security attributes, but the above instances provide a good minimum common denominator, a baseline of sorts. In 2004, the National Institute of Standards and Technology in US came up with a set of high level recommendations for design, development and operation of secure information systems, outlining 33 engineering principles for information technology security. We wrap this section with an extremely abridged summary of the recommendations next, and refer the readers to the original NIST documentation [30] for full details.

NIST's Engineering Principles Recommendations: The NIST recommendations view an information system through a life-

[29] September 30, 1999 CNN article 'Metric mishap caused loss of NASA orbiter': http://edition.cnn.com/TECH/space/9909/30/mars.metric.02/

[30] Stoneburner et al., 2004

cycle which is divided into five phases: (i)Č Initiation Phase, (ii) Development/Acquisition Phase, (iii) Implementation Phase, (iv) Operation/Maintenance Phase and (v) Disposal Phase. Each of the proposed principles are related to the different phases to a different extent (or not at all). Finally, the recommendations are grouped based on their relevance to different aspects of security - foundations, risk based, usability, improving resilience, reducing vulnerabilities and the central role of network and distribution in modern day information systems.

Concluding remarks

When it comes to cybersecurity, the attacker has asymmetric advantages.

- Siting duck syndrome: The target is essentially fixed, while the attacker(s)'s location or timing of attack is not.

- Moving goal post: Security measures need to be continuously upgraded to meet as new vulnerabilities and threats come to the fore, and a setup which is good at a given time point may be inadequate at a future time point.

- It's a one way street: It is easy to launch attacks, including using proxies which themselves are compromised, while it can be very difficult to identify, let alone counter-attack the source(s).

- Defend a boulevard of broken barriers: An attacker needs to succeed only once, while the defender needs to succeed in warding off repeated attacks, including many that follow unanticipated path.

- Heads, they win, tails, you lose: Even if the defender 'wins' in preventing a successful attack, the costs and damages could be significant, while for the attacker the cost is low, returns can be high, and it is very easy to decide to disengage at any moment.

- System complexity is an anathema to security: If the system itself is complex, then the chances of the existence of undiscovered (by the designers and operators) bugs and resulting vulnerability that can be exploited by an adversary increase. If the security controls are complex, chances of them not being used, or misconfigured, and so on, increase. All to the advantage of potential attackers.

- The ultimate bottomline is, there is no perfect security!

Since security is never going to be perfect, preventive measures need to be complemented with adequate mechanisms to detect security breaches, and respond to them, including having contingency plans. Furthermore, a defense in depth approach should be adopted to achieve layered security - deploying multiple (independent) mechanisms to defend against any specific vulnerability, so that even if some of the controls fail, there are further means to defend. Security should also not be solely reliant on technology, and requires proper processes and (correctly trained) people to make the best use of the technology at hand. Depending on the kind of activity an organization is involved in, there may also be laws and regulations that need to be adhered to. Likewise, industry standards form a minimal common denominator that all players of the specific industry are expected to adhere to. Ultimately, system security is an enabler for most, but not their bread and butter activity. This, along with the knowledge that security can never be absolute, implies that there will be a perennial tension between the amount of resources to allocate for security, versus the benefits derived therefrom. This requires a proper risk analysis to strike a balance and utilize resources judiciously. The rest of this course introduces the basics of all these issues — and is meant to provide a broad if somewhat shallow understanding of the nuances of security management — and should benefit a wide spectrum of audience with varied career aspirations: be it engineering resilient systems, information system administration, security consultant or (chief) information or information security officer of an organization.

Case study: Target targeted through HVAC

Sources:
http://krebsonsecurity.com/2014/02/email-attack-on-vendor-set-up-breach-at-target/
http://krebsonsecurity.com/2014/02/target-hackers-broke-in-via-hvac-company/
http://www.bloomberg.com/bw/articles/2014-03-13/target-missed-alarms-in-epic-hack-of-credit-card-data
http://www.ibtimes.com/timeline-targets-data-breach-aftermath-how-cybertheft-snowballed-giant-retailer
http://www.cio.com/article/2600345/security0/11-steps-attackers-took-to-crack-target.html

On 19th December 2013, US retailing giant Target corporation made a public acknowledgement that customer data had been breached. It will eventually turn out that more than a 100 million customers' personal and credit card data was stolen in this incident. The subsequent analysis and disclosures on how the hacking took place makes an interesting read, particularly given that the hack was not technologically speaking very sophisticated or unconventional, the necessary technology to prevent it was well in place, as was the time to react and prevent the attack, but nevertheless a series of missteps and inaction led to what is one of the biggest data breach incidents in the history. While a precise detail of how the attack happened is probably hard to find in the public domain, a vast wealth of information is nevertheless available, and provides ample insights for the purposes of this course.

The attack

It all started with a standard email phishing based hack of a heating, ventilation and air conditioning (HVAC) firm called Fazio Mechanical, infecting their network with a password stealing bot. While the HVAC firm did have a malware detection software running, it happened to be a free version which only supported on-demand scan, rather than real time monitoring. It also so happened, that the free version was meant for personal use, and not for corporate use, but that's somewhat besides the point but for to emphasize that the HVAC company did not consider securing its IT infrastructure seriously enough. Consequently, the compromise went undetected.

In the immediate aftermath of the attack, it was not clear why Target would have given an HVAC company external network access. Often HVAC's have access to their client's network in order to monitor and control the cooling and ventilation for purposes such as power savings and maintenance. But Fazio Mechanical's own words in their defense 'data connection with Target was exclusively for electronic billing, contract submission and project management'. Further puzzling question would be why Target?s payment system network would not be isolated from the same.

Subsequent dissection of the incident reveals the following trajectory that the attack followed. Using stolen credentials, the attackers could access Target's web application dedicated to vendors hosted in Target's internal network. The application however did not allow arbitrary command executions, as would be necessary to compromise the machine. However, its speculated that the attackers managed to upload a PHP file, possibly exploiting the fact that the application allowed file uploads (for say, invoices), but did not check if the uploaded file was an executable. The PHP file was named 'xmlrpc.php', a popular PHP component, thus hiding in plain sight a malicious code. Once a first piece of malicious code is injected in the system, it can be used as a backdoor to introduce further malicious code and execute arbitrary system commands.

At this stage of the attack, having established the capability to execute arbitrary system commands, the attackers had to carry out reconnaissance of the network to identify further vulnerabilities and locate digital assets worth stealing. Target's Active Directory containing the data on all members of the domain: users, computers and services was queried with internal Windows tools using the standard LDAP (Lightweight Directory Access Protocol) protocol. This might also have been aided by publicly available and easy to find (apparently benign) documents on Target's supplier portal, which contained embedded meta-information that revealed Target internal network information such as usernames and network layout. Once the target names were located, simple DNS query exposed the corresponding IP addresses.

The attack (continued)

Having identified the targets, access privilege for the same was the next objective. Access token from domain admins were stolen next, likely by using a rather standard technique called 'Pass-the-Hash' (Pass-the-Hash is a hacking technique that allows an attacker to authenticate to a remote server/service by replaying the underlying hash of a user's password, instead of requiring the associated plaintext password) to impersonate an Active Directory administrator, and use this domain admin privilege to create new domain admin accounts. Creating a new domain admin account with all necessary privileges isolated the attackers from being affected if and when the genuine administrators change password or witness any suspicious activity on their own accounts.

Equipped with the new credentials with administrative privilege, the attackers could access and propagate their malware to the target machines. There were a few minor obstacles like bypassing firewall, etc. which were however easily achieved at this juncture with the degree of control and privileges at their disposal. The attackers could now gain access to around 70 million customer records with personally identifiable information. However, Target, in compliance with Payment Card Industry Data Security Standard (PCI DSS), had not stored customer's credit card information.

Thus the attackers next proceeded to install custom made malware on the PoS (point of stale) systems to scan the memory of infected machines and save any credit cards found in a local file. It is speculated that an initial infection of PoS machines was carried out during the period of November 15-28 (Thanksgiving and the day before Black Friday), when the attackers could validate that the malware was working as desired, before deploying it across a much larger number of PoS machines within the next two days, to set up the infrastructure to collect credit card information from live transactions. Subsequently, an estimated 40 million credit card information would be stolen over the period of November 27 till December 15, 2013.

At this juncture at the end of November 2013, though the attackers had set up access to the sensitive PII and credit card data, it was being stored in a server within Target's own network, though under the attacker's control. They still needed a way to exfiltrate the data out, which they eventually orchestrated using multiple compromised machines across US as transitory drops, the digital footprint eventually going to Moscow.

It will ultimately be Federal law enforcement officials who would contact Target on 12th December 2013 to inform them about the breach, leading to eventual mitigating actions on Target's part to conform and eradicate the malware and stop the attack.

The victims, the silent accomplice

Though both the HVAC firm Fazio Mechanical as well as Target corporation are eventually victims of the said cybercrime, they are not totally blameless victims. We mentioned already the shortcomings at the HVAC's end. However, a lion's share of the blames lie with Target itself.

Months before the incident, in September 2013, Target was certified to have met the PCI standard. In the retail sector, Target was a leader in passing the certification. Roughly half an year prior to the incident, Target had also deployed a malware detection tool by a company called FireEye, who had among its customers even organizations like the Pentagon. A team in Bangalore monitored the Target computers round the clock, and were to notify Target's security operations center (SOC) in Minneapolis if anything suspicious is detected. So, at a first glance, it would appear that Target had done its due diligence.

On November 30, the hackers were ready with capturing the live transaction data, and effectuated the last phase of their attack - namely, setup further malware to move the stolen information out of the Target network to transitory servers spread across US to cover their tracks. At this juncture, the FireEye malware detection software raised alarm regarding suspicious activities, tagging it as the most urgent in FireEye's grade scale, which the team in Bangalore in turn alerted the Minneapolis SOC about. However, the alert was simply ignored - possibly (wrongly) considered as a false positive. In fact, subsequent analysis of logs will show that there was a further alarm on 2nd December, when the hackers had upgraded their malware, but the SOC had again ignored it.

The victims (continued)

The irony is, that the system could even have been configured to automatically eliminate malware, to obviate human intervention. However, this had been deliberately turned off by the security administrators. Given that it was still a relatively newly introduced technology at Target, it was being tested and not fully trusted, and thus it was configured not to make autonomous decisions, and likewise, its alarms were not taken diligently.

Why do we fall?

In every security incident, there are lessons to learn, some new ones, some common knowledge among security experts but nevertheless easily ignored amidst day to day grind, and thus worth reemphasis.
Meeting an industry (PCI) standard by no means mean that one is secure. One should understand that it takes years of work to draft and ratify a standard, but in the meanwhile the security landscape changes significantly. A standard ought to be viewed as a minimal denominator — the least one should be doing — and it by no means demonstrate best effort. On the contrary, being certified may make one complicit, or grasp for excuses to shift blame in the event of a security incident.
Technology in itself is inadequate in addressing all our security needs. Foremost, there may not even be adequate technological remedy to all our security problems, but even when the technological means do exist — as was the case in this particular incident — the organization's processes (or lack there of), and the people in the loop could become bottlenecks. Be it configuring the FireWire software poorly, or failing to respond upon detection by ignoring the alarms from FireEye, or in failing to detect new domain admin accounts or monitor access patterns of these newly created accounts, or, to start with, develop a web application with the correct preventive measures that would have disallowed the uploading of an executable file through the web application and nipped the attack in its bud, the failures were all-rounded across the spectrum of prevention, detection and response, and the people and processes in place are largely to be blamed for these. Lack of proper isolation (not a well set up DMZ) for the part of network accessible to third party vendors from the internal networks and the payment network also indicate some poor system setup choices.
The consequences of the breach would also have been substantially less, if an added layer of security, as afforded by Chip and PIN, had been deployed (as is widely the case in Europe). This again shows how, even if technological solutions may exist, inertia may dictate poorer security.
Ultimately, preventive measures will inevitably fall short and attackers would find some way in or other. However, lack of (or, in this case, ignoring) real time detection, and failing to respond fast are ill afforded.

Epilogue

As an aftermath of the incident, Target lost consumer confidence and sales, the consequent financial downturn led to job cuts of many ordinary workers at Target. The presiding CIO and CEO also had to step down in March and May 2014 respectively. Job cuts continue at Target, but probably these are effects of multiple factors and not directly consequent of only the data breach anymore. FireEye's share prices sky rocketed after details of the breach came to fore, though the prices as of August 2015 are significantly lower than the peak it had reached in February 2014. In February 2015, Target revealed its estimated direct financial loss from the data breach to be around 162 million dollars, though such estimates are likely very coarse. In March 2015, Target also agreed to pay 10 million dollars as a settlement to a class-action lawsuit, whereby customers who had been affected by the breach (if they can prove so) could get compensated by up to US$ 10,000. Analysis of the malware code revealed some information indicating possible involvement of Rescator.so (a popular) in the data breach, which was also involved in selling the stolen data. Apparently Rescator.so itself was hacked in March 2014 and logins, passwords, and payment information of carders (people who buy and then use stolen credit card information) was posted online. Whether this poetic justice was delivered by any of Rescator.so's victims or one of their business competitors is anybody's guess.

Bibliography

Europol. EU Organized Crime Threat Assessment. OCTA 2011, 2011a.

Europol. Major international network of payment card fraudsters dismantled, 2011b. URL https://www.europol.europa.eu/content/press/major-international-network-payment-card-fraudsters-dismantled-1001.

David Wolman. *The End of Money: Counterfeiters, Preachers, Techies, Dreamers–and the Coming Cashless Society*. Da Capo Press, 2012.

Christopher Hadnagy. *Social Engineering: The Art of Human Hacking*. Wiley Publishing, 2011.

Charles P. Pfleeger and Shari Lawrence Pfleeger. *Security in Computing*. Pearson, 2007.

Rebecca Boyle. Hackers could access pacemakers from a distance and deliver deadly shocks, 2012. URL http://www.popsci.com/technology/article/2012-10/hacker-attackers-could-reverse-pacemakers-distance-delivering-deadly-shocks.

Alan Grau. Hackers invade hospital networks through insecure medical equipment. *IEEE Spectrum*, 2015.

John D. Howard and Thomas A. Longstaff. A common language for computer security incidents. Sandia report, 1998.

ISO7498-2:1989. Information processing systems – open systems interconnection – basic reference model – part 2: Security architecture. Technical Report ISO7498-2:1989, International Organization for Standardization (ISO), 1989.

Gary Stoneburner, Clark Hayden, , and Alexis Feringa. Engineering principles for information technology security (a baseline for achieving security), revision a. Technical Report NIST Special Publication 800-27 Rev A, National Institute of Standards and Technology, 2004.