Computer Security Impaired by Legitimate Users

Denis Besnard & Budi Arief

denis.besnard@ncl.ac.uk
l.b.arief@ncl.ac.uk

Centre for Software Reliability
School of Computing Science
University of Newcastle upon Tyne
Newcastle upon Tyne NE1 7RU
United Kingdom

Abstract. Computer security has traditionally been assessed from a technical point of view. One other view is about the role played by legitimate users of systems in impairing the level of protection. In order to address this issue, we wish to adopt a multidisciplinary standpoint and investigate some of the human aspects involved in computer security. From research in psychology, it is known that people make biased decisions. They sometimes overlook rules in order to gain maximum benefits for the cost of a given action. This situation leads to insidious security lapses whereby the level of protection is traded-off against usability. In this paper, we highlight the cognitive processes underlying such security impairments. At the end of the paper, we propose a short usability-centered set of recommendations.

Keywords. Computer security, cognitive psychology, cost-benefit trade-offs, work practices, risk

1.   INTRODUCTION

Our society is becoming more and more dependent on computer systems, which nowadays are used in everyday life, from business to banking, from entertainment to healthcare. Most of these systems are interconnected through the internet, which inherently is very open and vulnerable to cyber-attacks. Attacks on these systems cause a wide variety of disruptions, ranging from losses in service to financial or safety consequences. The threats that these attacks constitute over the mere use and survival of IT systems have led to wide quantitative surveys (e.g. [20, 24]).

Because a single attack can freeze an entire sector within hours, securing computer systems has become a very important part of system design, development and deployment. However, the measures implemented are not always as efficient as required. More research and better tools are therefore needed in order to understand, block and anticipate security threats. A novelty may be the growing acknowledgement from the research community that technical solutions alone are not enough any more. Historically, security has been very often addressed from the attackers' side. From this angle, the emphasis has classically been on the motivations and means used to break into systems [22, 32, 38, 42, 56]. However productive this research area has been and still is, it tends to blur the exact role of the legitimate users (e.g. end-users, security officers, managers, designers) who are also actively involved in computer security impairments. It seems to the authors that this other angle is worth exploring as well. Moreover, as part of the DIRC[1] research,

---

[1] DIRC (Dependability: an Interdisciplinary Research Collaboration) is a UK-based interdisciplinary research project on the dependability of computer-based systems. Visit DIRC at http://www.dirc.org.uk.

we are interested in interdisciplinary aspects of computing. For these reasons, some aspects of the role played by legitimate users with regard to security are addressed in this paper.

Computer security is an area that cognitive scientists have not investigated as deeply as human-computer interaction or problem solving. It nonetheless offers an interesting aspect in the sense that there are conflicting objectives held by some of the actors of a single system, namely attackers and legitimate users. It follows that depending on the goal that an actor is pursuing (attack or legal use), the use of a given computer system will differ dramatically. Whereas the roles of attackers are pretty clear (e.g. intrusion, denial of service), those of legitimate users regarding security are more subtle. Stemming from this assumption, we will examine some of the latter's practices and shed some light on the mental processes involved. We will try to assess the extent to which computer security can be interpreted in terms of an intuitive cost-benefit trade-off.

The rest of the paper is laid out as follows. The next section (section 2) provides some background about trade-offs in the workplace. We then investigate the cognitive aspects of trade-offs (section 3) and expose some examples of insecure computing practices (section 4). The paper then discusses risk issues, the antagonism among some of the security actors, as well as organisational policies (section 5). This is followed by a set of recommendations (section 6), bearing in mind some limitations (section 7). The article concludes with some brief reflections on the status of computing in modern society.

2.   TRADE-OFFS IN THE WILD

Since Simon [55] and his concept of *bounded rationality*, it is accepted that human actions do not reach perfection but instead seek an acceptable level of performance with respect to their goals and what the cognitive resources allow. The fact that the cognitive system never aims at handling all the data available in the environment is a central aspect of the cognitive resources saving strategy. As a consequence, cognitive acts are an intuitive and implicit trade-off which balances cost and efficiency [10, 57]. This strategy is put in place for the majority of human actions and therefore applies to an extremely wide class of situations, e.g. troubleshooting [13], medical prescriptions [25], control of dynamic situations [7]. Just like in other activities, trade-offs introduce a risk (see for instance [36]) by not taking into account some possible consequences deriving from the decided actions. In the case of IT security, trade-offs and risk can thus combine themselves and implement a threat. Therefore, our purpose will be to assess the relevance of trade-off mechanisms regarding security impairments.

Before we consider computer security, it can be useful to have a look at a field example. Although the latter is quite remote from computing, we think it puts things clearly and shows how the parameters of a trade-off are handled by humans.

On December 30, 1999, in Tokaimura (Japan), an accident occurred at the JCO nuclear fuel processing plant, causing the death of two workers (see [29]). The immediate cause of the accident was the pouring of approximately 15kg of uranium into a precipitation tank. Since there is a limited amount of uranium that can be put together without initiating fission, this procedure required mass and volume control. As the critical mass was exceeded, a chain reaction occurred, generating lethal radiations. The workers' task was to process seven batches of uranium in order to produce a uranium solution. The tank required to process this solution is called a *buffer column*. At JCO, its dimensions were 17.5 cm in diameter and 2.2 m in depth, a geometry permitting a better control of fission reactions. The inside of this tank was known to be difficult to clean. In addition, it was positioned only 10 cm above the floor, making it difficult to collect the uranium solution from the bottom of the column. Thus, workers illegally opted for using another (larger) tank called *precipitation tank*. Due to its dimensions, this latter tank was not geometrically safe but

it was positioned 1 m above the floor. Moreover, it was equipped with a stir propeller making it easier to use for homogenising the uranium solution. The pouring of the 15kg of uranium at once triggered the criticality accident. Its causes were rooted in a complex combination of deviant organisational practices. Among these, pressures from the managerial team to increase the production without enough regard to safety implications and lack of crew training played a significant role.

In hindsight, we speculate that the operators have traded-off productivity and practicality against risk. As their knowledge about critical uranium masses was poor, they were unaware that they were crossing a safety boundary. This case is an instance of how trade-offs can go wrong. With this example, we want to highlight the workarounds that operators often implement in order to perform daily actions in a less constrained manner (see [30]). These workarounds can be put in place in a wild way, and depending on the level of knowledge and perceived risk, getting the work done sometimes overrides security concerns. Also, we wish to highlight the role of the managerial team at JCO who played an active goal in the triggering of the accident [29]. This latter point is of relevance to our paper and will be addressed later.

Following a cognitive approach, we believe that virtually every decision is a matter of trade-off. Humans do not try to produce perfect responses to the environment. Instead, they tend to accept good enough solutions. We think that this conception of human cognitive activities applies in computer security, for both attackers and legitimate users. The former attempt to design effective worms or denials of service, for instance. The latter, in turn, try to protect themselves as effectively as possible. But in both cases, there are not infinite amounts of resources (e.g. time, money or effort) to allocate to attacking or protecting. This is where human flexibility comes into play: people perform intuitive trade-offs between (some form of) cost and (some form of) benefits. We will consider some concrete computer examples after having briefly explained how the concept of trade-off translates in security.

3. TRADE-OFFS FROM A COGNITIVE PERSPECTIVE

To better illustrate how we use the concept of trade-off, we represent it graphically in Figure 1. The dark area at the lower right-hand corner represents the maximum efficiency where one reaches high benefits for low costs. The top left-hand corner, on the contrary, represents a poor efficiency where one spends a lot to gain little. Between these two extremities, there is clearly an entire continuum. The theoretical trade-off line represents a frontier between costs and benefits. Any activity above this line will cost more than it rewards. Conversely, any activity below this line will reward more than it costs.

Having said this, humans do not always obey logic but seek cheap actions with maximum expected benefits. Therefore, the least-effort trade-off line may be a more realistic one. This view is derived from a *least effort* rule whereby humans attempt to reach an acceptable level of performance with the minimal mental effort [11]. As a consequence, decision making can become a biased benefits-driven process. When applied to a security-usability trade-off, usability may come first, hence turning security into a side-issue. This will be further discussed in section 5.2.

Let us now apply the graph to the simple example of a user considering adding a button to a toolbar in order to access a function more quickly. If the button is going to be used only one time, the time cost of adding it may be higher than the expected time saved during the task. Therefore, it is likely that the button will not be added. On the other hand, if this function is to be used repeatedly, the expected benefits may be worth the time spent in configuring the system. With this simple example, we wish to recall that humans intuitively, though implicitly, evaluate the efficiency of their decisions before they implement them.
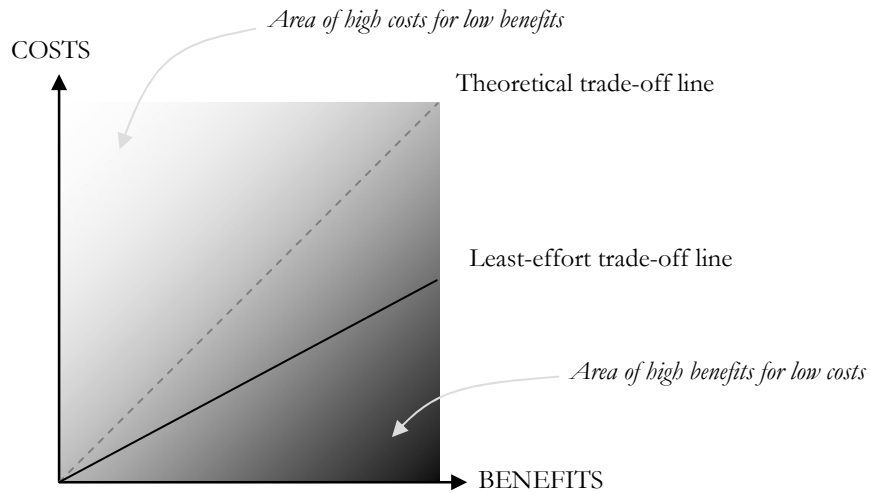
Figure 1: Graphical representation of a costs/benefits trade-off.

There are cases where human actions are given explicit limits: decisions can be benefit-driven or cost-driven (see Figure 2). In the first case, a course of actions is interrupted when some objective is met. In the other case, the target is set in terms of cost (money, time, etc.) and actions will stop when the limit is reached. In both cases, the course of actions never follows the straight trade-off line. Instead, we believe it fluctuates with time depending on the given phases of the work. For instance, one may be prepared to temporarily carry out a costly action if high benefits are expected from it later.
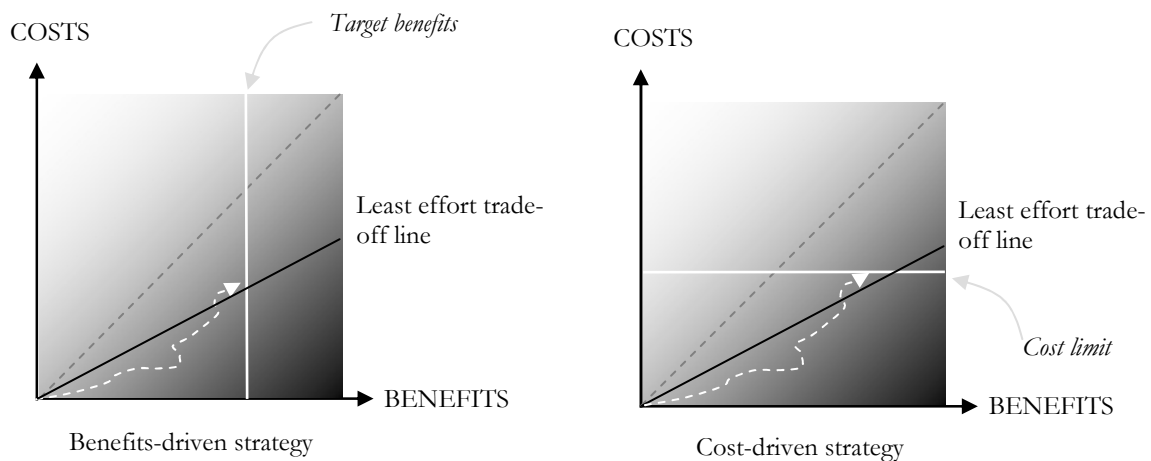

Figure 2: Graphical representations of benefits-driven (left) and costs-driven (right) strategies.

Classically, attackers are said to exploit security holes left open because of a poor design and/or insecure practices. In other words, the malicious intentions of the attackers are, to some extent, facilitated by the behaviour of some legitimate users. We obviously do not put the blame on them. Neither do we believe that the motives that some attackers promote (e.g. learning, curiosity, challenge, etc.) will ever justify any sort of damage caused to someone's data, tool or service. Having said that, security is a two-way issue. Merely assessing it from the attacker's point

of view only captures half of the problem (see [8] for a survey on attackers' techniques and motivations). The other half is about how we (legitimate users) use our computers.

4. SECURITY TRADE-OFFS BY LEGITIMATE USERS

From our cognitive standpoint, administrators as well as end-users consider their actions from an economic point of view and trade-off security against usability. This practice, as highlighted in the following sub-sections, introduces threats in systems that attackers exploit. It is not the case that attacks can be easily eliminated, but understanding where trade-offs lie may allow system designers to think more about the interest of making security products and policies compatible with some intuitive notion of usability. This is not only an issue about "comfort of use" but, as we will see in the following sub-sections, is a problem that directly impacts security.

4.1.1. Passwords: a memory issue

Although new approaches towards authentication have been proposed [15, 17, 39], passwords still remain a widespread security mechanism. A number of modern software products still tend to force passwords of minimum eight characters long. This tendency probably originates from a desire to control accesses more tightly with the hidden assumption that it will increase privacy of data. Although they are harder to crack, long passwords are not totally secure [46]. When one actually looks at what happens at the workplace, human cognitive limitations become obvious: users cannot remember their passwords and need external memories (e.g. sticky notes on monitors). The use of passwords raises several usability problems [6]. Security faces a nice paradox where by increasing the complexity and number of passwords, the level of protection could actually decrease [59].

User login and passwords to computer accounts are used very often and are the main method to get access to systems. In this case, frequency compensates for complexity: the password is used often enough to be remembered [54]. But how many counter-examples are there where people have to write down passwords? To cope with this problem, some systems (most notably e-commerce web sites) offer to remember them. Again, for the sake of usability, a user may be tempted to use such storage features (cookies). It is a useful feature but it comes down to the user's judgement as to whether a service or a piece of data is trivial enough so that its password can be stored on a computer. Another drawback is that users may not remember their password if somehow the cookies are emptied (e.g. when the system is rebuilt).

The same kind of problem applies to BIOS (Basic Input/Output System) password. Administrators sometimes keep the default password, information that can then be found by anyone on the internet (see for instance [1]).

4.1.2. Anti-virus software updates: a risk issue

Anti-virus protections are useful barriers but only when they are up-to-date. They need some attention in this respect. But maintaining, updating and upgrading them has a cost that can conflict with end-user's main task, thereby impacting security. Therefore anti-virus protections potentially leave a hole open, the size of which depending, among others, on the frequency of updates and reactivity to patches availability [51]. Hence risk, which is typically perceived inaccurately by humans [50], comes into the equation at this level. End-users or security officers have to accept a certain level of (perceived) vulnerability regarding their system. Automatic updates have been felt to tackle this problem by lowering the likelihood of holes in the anti-virus protection. It seems to be a reasonable belief and is indeed a widespread feature. Unfortunately, the update process can be corrupted by an attacker. Although this is not a trivial task, it has been reported that centralised distribution of software could be tampered [2, 31]. This puts high threats on IT systems since it can offer an attacker to automate the installation of backdoors or the downloading of harmful contents on the end-user's machine. For this reason, automatic

updates are not a panacea. Despite this negative state of facts, large organisations use this feature extensively. They may assume quite reasonably that the gain in usability and the regularity of updates together will provide benefits that are felt to overweigh risk.

4.1.3. *Email attachments: a trust issue*

Email attachments have been used very widely for spreading malicious code [22, 63]. Typically, the code is added as an attachment to a seemingly benign email. When the attachment is opened, the code is executed on the machine, exploiting security holes in the email program. Harmful email attachments pose two problems. First, they very often come from trusted senders who were infected themselves. Thus, due to this trust relationship, the degree of suspicion regarding the decision to open the document is already low. Second, email attachments are used so widely for legitimate purposes that opening them has become as automatic as picking up the phone [9]. Of course, a rule such as "*do not open attachments you are not expecting or from people you do not know*" is unworkable since it generates too many false positives: many valid attachments would have to be left unread. Again, automation has been felt to tackle this problem: many mail servers now include anti-virus scanners. However, this is not the optimum solution since these scanners a) may not be up-to-date, b) sometimes detect false positives and c) open a discussion on whether or not scanning might invade privacy.

4.1.4. *Files sharing: a practicality issue*

On Windows operating systems, there is a feature that allows users to share files on their computer with other users. This is a very useful feature, especially for collaborative work among several users. It facilitates an easy and quick way to access or modify (potentially big) files from multiple computers. On the down side, this service (running on port 445) is vulnerable to denial of service attack [43] or to worms such as W32/Deloder [45].

For everyday users, the practicality of sharing file of folders may outweigh security concerns. As in most cases, the threats can be minimised by applying security patches. However, the existence of such patches is not known to everyone. Moreover, the installation of the patch itself consumes work time and causes a distraction, which can be perceived altogether as factors of resistance to securing the system. Lastly, users may simply not be aware of the risks associated with the default settings in file sharing and they may "*share with everyone*" without knowing how broadly they disclose their data.

5. DISCUSSION

After having considered some concrete examples based on the use of computers, it seems necessary to adopt a broader view and address some more general issues. Among these, accepted losses and risks have to be mentioned as security policies are not meant to protect every single piece of data or service. Furthermore, threats can be discussed from the standpoint of an antagonism between the roles of attackers and legitimate users. The discrepancy in the motives of these actors is where security holes lie. Lastly, we will quickly evoke organisational issues by describing a multi-layered view of systems' security.

5.1. Accepted losses, risk perception and systems' protection

Although we have reasoned so far under the reductive assumption that all the data have to be protected, we now want to highlight a somewhat different picture according to which there exist some acceptable losses that humans implicitly take into account in setting up protections. The trade-off here involves the cost of protection and the cost of loss (see [26]). Data or services that can be easily replaced, disclosed or lost without serious consequences will probably have a relatively low level of protection. The underlying evaluation of the required level of protection is believed to be done in an intuitive manner most of the time. We also think it guides, to some

extent, the security policies adopted by organisations. As each single piece of data cannot be equally protected, some of these data are inevitably left vulnerable to attacks. This may be a sensible decision if, as aforementioned, loss or disclosure of data is accepted. In this view, it seems important to highlight that protecting a system is an implicit dialogue between security officers and attackers. We believe security policies define the nature of this dialogue before the occurrence of any attack. To some extent, this conception goes against the widespread belief according to which "attackers play first".

The following point may be perceived as a side-issue in this paper but the information that is intentionally left unprotected can cause indirect damages. A category of attackers called *social engineers* are specialised in gathering this kind of data by the means of psychological manipulation. They seek contact with a legitimate insider of the target system and trick that person into revealing passwords or other information that compromises the target system's security. The attackers usually conduct their deception through the phone in order to minimise the risk of being caught or recognised. See [34, 35, 44] for more information on social engineering.

The cost of losing valuable data or service may be one of the obvious drivers for designing and applying security protections. But benign or seemingly trivial pieces of information, although not directly security-critical, can be damaging as well, especially at the hand of malicious social engineers.

IT security shares some similarities with many domains. Risk taking is one of them and its management can be both unavoidable and tacit (see [25] about general practitioners prescriptions). Let us take the following example. A car is safer when it is immobile than when one is driving it. But for a car to deliver its service (transporting people and goods), the driver and the passengers are forced to expose themselves to risks. These may be reduced down to some acceptable level if the driver is careful and experienced. But there will always be a number of factors he or she will not be able to influence (other drivers, mechanical incidents, etc.) that will impact the level of safety of this situation. The same argument holds with, for instance, a server. Not plugging it into a network is a relatively secure condition but the service will not be delivered. Therefore, some risk has to be accepted for virtually any piece of equipment to fulfil its function.

Another problematic and risky situation is one where the benefits expected from some programming decision outweigh protection measures (see [49], about programming with COTS). Saving time and cost is common practice within the community of programmers [52] and may lead to biased benefits-driven decisions [18]. This applies to security in the sense that a threat can be identified by e.g. a software developer but found too costly to fix or thought unlikely to be exploited. This behaviour extends to an extremely wide range of cases. For example, in everyday's life, we tend not to wear our safety belt when driving on very short distances (e.g. parking our car in the driveway). In this situation, we perceive risks as being very low and we implicitly adapt our protection level accordingly. The smaller the perceived risk, the lower the level of protection. However, humans are typically biased at perceiving *actual* levels of risk [50] and rarely have an exhaustive knowledge of the systems they interact with. It follows that the impact of a given practice over the security of a system is unlikely to be accurately assessed by an end-user. Thus intuitive, heuristic risk assessments, although acceptable for everyday's life, do not always accurately capture the criticality of certain threats. This inaccuracy therefore degrades the identification and compensation of security breaches, which in the end can depend on subjective decisions.

Last but not least, there is a human tendency to "slide on the risk slope". Large security incidents or industrial accidents are not caused by a sudden change in security or safety policy. Departure from a reasonable level of risk does not happen in one day. It is an accumulation of a number of

small insecure increments that progressively deteriorate the level of protection, each of which being seen as acceptable *per se*. This is a classic situation in large industrial system's safety: large-scale accidents are made of a concatenation of small failures [41].

5.2. Antagonism among security actors

It seems plausible that attackers, just as legitimate users do, perform trade-offs in the way they use their own computers. They may tend to intuitively and implicitly compare the costs of their actions (e.g. time) to the expected benefits and then take decisions on the basis of this evaluation. The rule-of-thumb states that if costs are perceived as worth the expected benefits, then some action is likely to be performed. However, because attacker's and legitimate users' motivations are fundamentally different, we think that their respective trade-offs are different in nature. Attackers attack because they get a reward of some sort (self-satisfaction, peer-recognition, money, etc.). Legitimate users protect themselves because they need to. As far as trade-offs are concerned, these motives bring a consequence that attackers may care less about costs than legitimate users do. This discrepancy of motivations is perhaps where threats originate.

In the case where they have identified a target, attackers may be more focussed on the damages expected from their attack than on the costs involved. On the other hand, legitimate users may prioritise usability with little concern about security. This is a common case of usability-driven behaviour. We represent this discrepancy as a gap between attackers' and legitimate users' trade-off strategies (Figure 3). In our conception, this gap gives some advantage to attackers. The larger the gap, the most successful the attack could be.
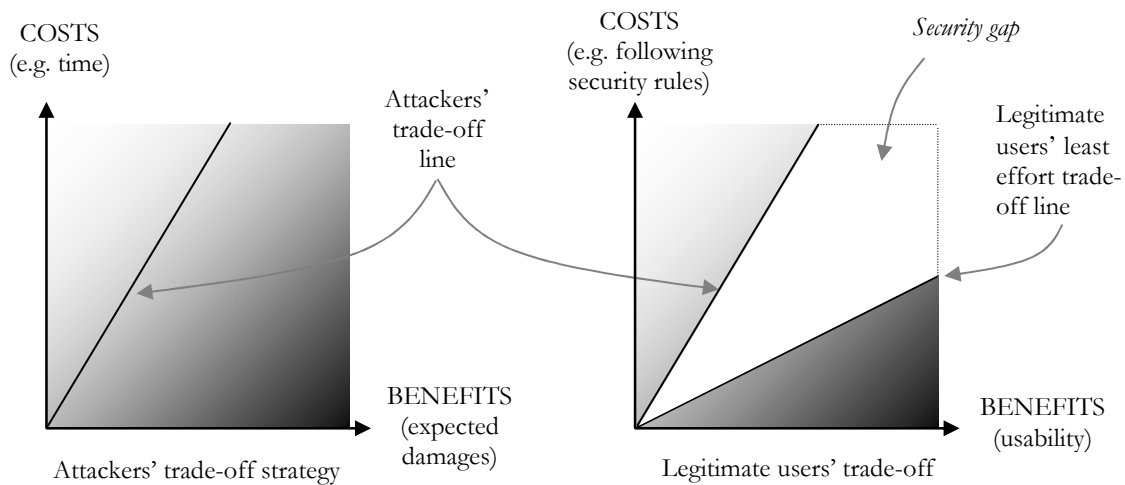

Figure 3: Graphical representations of two different trade-off strategies leading to a security gap.

From our point of view, attackers' trade-off involves expected damages against cost of actions whereas legitimate users trade-off usability against security. Having said this, it is not the case that an action making a system more usable will systematically degrade security. Instead, the point is that legitimate users, sometimes without being aware of it, prioritise usability at the detriment of security. In our opinion, this creates a security gap. Therefore, a successful attack can be described in terms of a malicious action whose degree of refinement is higher than the degree of protection of the target system.

5.3. Beyond the individual picture…

So far, we have been concentrating on an individual perspective where cognitive factors are thought to play a determinant role. Beyond this picture, we want to acknowledge the collective dimension of security in large distributed computer-based systems. More precisely, we think that Reason's model [48] (see Figure 4) adopts a useful view on organisations in the sense that they are described in terms of multi-layered systems. Applied to the field of security, this view can help describe a computer-based system as one composed of threats, actors and protection layers. With this model, security is described as a multi-layered process where a variety of users (e.g. developers, security officers, end-users…) have a role to play. Each of these users impact on security. Administrators and/or end-users, for instance, by not making updates for their anti-virus software, leave holes open for attacks. This type of local failure may exist at any given layer of the organisation, for any role. It creates latent security breaches that, combined with each other, can defeat an entire system's protections. In the context of this paper, these breaches are interpreted in terms of trade-offs whereby users simply wish to reach good enough solutions.

When applied to security, Reason's model can describe, from a system point of view, the impairments made to security by legitimate users and attackers. The latter attempt to propagate attacks through security holes in order to reach an objective such as data, a service or to cause disruption in the functioning of the system.


Figure 4: Successful attacks propagate through several protection layers (adapted from [16]).

How people perceive IT security [23, 27] and privacy [4, 19] is a useful approach to assess individual contributions or impairments to security. However, these factors have to be brought back into a broader picture where security breaches are caused, facilitated or maintained by the combination of a variety of causes rather than by mere, isolated end-users' actions. By quoting Reason's model, our intention is to highlight the combination of factors needed for an attack to succeed. According to the popular belief, attacks occur because some malicious people exploit security holes. We wish to promote a somewhat different view according to which successful attacks are a combination of weak protections and malicious intentions. This idea can be stretched even further. As we cannot eliminate attacks altogether, the most productive approach may be to regard attacks as the outcomes of flawed policies and/or practices. Their origins are

deeply rooted within early design assumptions or within managerial decisions. For instance, due to productivity constraints, the manager of a small company may misjudge the importance of protecting IT activities. This can take the form of a backlog of security actions waiting to be done. Such a *laissez-faire* policy may lead to the adoption of a poor security culture. The latter may propagate through the various stakeholders of the organisation, leading to e.g. unprotected data or weak passwords. There could obviously be an infinite number of examples that would follow the same pattern of a multiple, intricate set of causes (see [16] for more complete views on Reason's model application to security).

5.4. Summary

Here are the points that we have laid out in this paper:

- Trade-offs are sometimes implemented in a wild, uncontrolled manner. Legitimate users sometimes prioritise immediate benefits to the detriment of long-term security.
- Passwords, anti-virus updates, email attachments and shared folders respectively raise such issues as memory limitations, risk, trust and practicality.
- Security does not imply protecting everything since some losses are acceptable. However, because risk perception by humans is highly biased, valuable data could be under threats.
- Trade-offs by legitimate users differ in nature from the ones performed by attackers. The resulting gap creates or maintains security breaches.
- Computer security is an organisational matter.

6. WHAT CAN WE DO?

Telling people what to do about security is one option. But one lesson that can be drawn from violations in systems is that one should not expect humans to always act as prescribed. Within industrial settings, procedures themselves do not rule the human behaviour [28] and there are many ways in which humans can configure a system and use it in unexpected and/or unprotected modes, even if it implies implementing a violation [5]. This seems to be a universal pattern and to this respect, at least, IT security is similar to virtually any other field of activity. The motivation for diverging from recommended practice may be based on an intuitive cost–benefit evaluation where potential negative consequences of one's act are overweighed by expected benefits. This is typical for passwords that are written down or passed on to colleagues. It is also true for harmful email attachments that happen to hit computer-aware staff in academic departments every now and then. Generally speaking, if the perceived risk attached to an illegal action (e.g. lending a password) is seen as lower than the expected benefits (e.g. gain in time), then a violation will be put in place. This is extremely common practice and goes well beyond computer security. In this trade-off, factors such as security culture and risk perception are key notions. And whether or not the user has a relevant knowledge of the potential consequences of his/her actions is what partly determines the level of risk involved and the final security of the system.

6.1. Recommendations

To put things simply, humans obey least-effort rules because they are cognitive machines that attempt to cheaply reach flexible objectives rather than to act perfectly towards fixed targets. As a consequence, each time an opportunity to do so arises, efforts are avoided. This rule applies even to the detriment of performance or security. From our point of view, this leads to some simple recommendations concerning both end-users and administrators.

- *Educate staff.* Education will not solve all the problems but will at least allow users to be aware of the consequences of their actions. Contrarily to safety at work which is regulated and trained for, IT security still seems to be poorly addressed (as reported in [14] and in [24]). This may be caused by the fact that most computers in organisations are used as clerical or editing tools that contain documents that are not felt to be security critical.

However, it is a very risky assumption to believe that because one is not holding or processing sensitive data, attacks do not have to be cared about. Nowadays, the reality is that being plugged into the internet is enough to be a target. Therefore, any member of staff interacting with computers should be aware of the damages caused by insecure practices. It is far from being the case. The reason may have to do with the intangible nature of electronic data flows, making most of IT security problems obvious only when attacks have succeeded. Because of the poor visibility of insecure settings, it is necessary to explain to staff that e.g. email attachments can be harmful, how intrusions are performed and how to choose and use passwords.

- *Procedures do not rule human behaviour.* Humans have an extreme ability for tweaking rules and procedures. Thus, IT administrators should not assume that staff members will follow rules to the letter. Instead, it seems more reasonable to assume that people will always find a way to do what they want, via a violation if needed. It does not mean that rules and procedures are useless. But just as obedience to the rules *per se* does not automatically increase safety at work [28], it does not increase security either. Instead, as Dekker [21] suggests, procedures should be seen as resources for action instead of an expectation about human behaviour. Procedures must be understood. Their efficiency relies more on the knowledge they require than on their blind acceptance.

- *Security must be user-centred* [64]. Generally speaking, the design of security products and policies should rely more on the rules of human-computer interaction, as suggested in [33, 47, 53]. Also, products should be designed in such a way that users can make sense of their properties [37, 62]. At a finer-grained level, passwords must be, at least, easy to remember and reduced in number as much as possible. As far as end-users are concerned, the ideal number of passwords is zero. It may seem an unworkable view to security officers but the reason why security policies have to be enforced to humans is because these policies require an effort from them. And rules that are felt too costly to follow are simply not respected [58]. On the usability front, studies revealed holes caused by security products that are difficult to use [60, 61]. Therefore, any measure getting closer to an *effortless* security policy is a step forward. People should not have to remember about IT security or even think about it. The entire workplace should be designed according to this principle. User agents as described by [3] are an instance of such an approach. Also, retina control and fingerprints, although not a panacea, are pure effortless authentication items. They are both unique and virtually impossible to share with others. Purists of IT security will object that they can be falsified or that they may not be usable in all physical environments (e.g. dusty, dark places, etc.). It is true, but how good is the situation *right now*?

- *Security is not end-users' task.* How secure a system is partly depends on how high security is set on the scale of an organisation's objectives. Security might be a relatively obvious goal for a system administrator. This it is not the case for an end-user. Solutions have to be thought of in order to make security transparent for whom it is not a primary objective. For instance, it should not be expected from staff members to report security problems, update virus protections or remember half a dozen passwords. These expectations are rarely met in the real world. Security is administrators' task. End-users should not be expected to always collaborate.

- *Be aware of contradictory objectives.* Asking staff to carry out their duty and spend time on updates and/or scanning files cannot be done at the same time and have to be traded-off against each other. Contradictory objectives, which are often inevitable, must be compensated for: security does not come first in end-users' mind (see [54]).

End-users will always have something else to do other than thinking about security. It seems to us that the idea of a *user-centred security* for end-users is a useful policy driver. Any measure going in this direction may improve systems' security.

7.   LIMITS

We have explained how, in our opinion, trade-offs between e.g. usability and security could impair the level of protection of a system. This surely accounts for the success of some intrusions. But this aspect of human functioning can be seen in a more positive way as not complying to the rules can also generate beneficial behaviours [12]. According to this view, violations are reconsidered under the angle of ad-hoc contributions to security, happening under exceptional circumstances and outside the frame of any clearly identified procedure. An example is unplugging the network cable from a connected machine when a suspicious behaviour is detected. It may not preclude any damage on this specific computer but it will prevent the attack to spread to other machines. This is the type of actions that designers probably do not expect users to take but that can nonetheless be implemented on-the-fly, thanks to human's intrinsic flexibility. This kind of unexpected contribution to security is hardly ever addressed in computer security but nevertheless deserves some attention.

Here is an angle that we have not considered in this paper. When attacks or intrusions fail, it may be that attackers faced problems that would have been too long to surmount given their level of competence, available time, or expected reward. Thus, the cognitive approach may also be fruitful for the study of attackers' failures. Obviously, the major challenge here is getting the data.

8.   CONCLUSION

Legitimate users who are not security-aware (e.g. researchers, clerical staff, managers) have other tasks to perform than spending their work time on securing their system. When the task of protecting a machine is felt to get in the way to the completion of their main tasks, users will probably overlook security if this allows them to ease their work. Harmful, usability-driven trade-offs are then put in place and create holes in systems' protections.

Understanding where trade-offs lie can allow a better understanding of the mental processes involved in security practices. In the case of legitimate users, we defended the idea that security is impaired because it is traded-off against usability or efficiency. Now, looking at the big picture, it may be that past engineering experiences have something to teach us. As Leveson [40] reports, the introduction of high-pressure boilers aboard steamers caused a myriad of accidents. Different people took different positions regarding this problem. Instead of banning the new boilers, it was suggested that risks could be limited by adopting simple designs. This way, a more careful implementation of new steam engines would leave time for scientific knowledge to build up. Ideally, this policy would lead to a safer technological evolution, based on another strategy than blind trial and error. It seems to the authors that the situation in IT security today is not so different from what it was in the steam engines era. We are trying to prevent security breaches within systems that sometimes exceed our level of understanding. There is an uncountable number of examples of this. For instance, all major software development companies are continuously issuing security patches for their products. The reason is that nobody can detect all the flaws in a large program. Nor can anyone foresee the creativity of malicious people in exploiting these flaws. The consequences are that a) the arsenal of currently available protection tools (anti-virus software, firewalls, access control, intrusion detection systems, etc.) are no guarantee against attacks and b) financial losses persist [20] and [24].

In the authors' opinion, the way computers are used by legitimate users accounts for a number of security breaches. Since there is no evidence that we will step back from the pervasive use of information technology in the near future, understanding and compensating individual insecure practices is still of immense interest. Until this happens, the security challenges that our society is

continuously facing and the financial costs involved will remind us that we have brought IT systems beyond end-users' control.

9.   ACKNOWLEDGEMENTS

10.   REFERENCES

[1]   "Default Password List", available online at http://www.phenoelit.de/dpl/dpl.html .

[2]   "Gnuftp Compromise", available online at ftp://ftp.gnu.org/MISSING-FILES.README (2003).

[3]   Ackerman, M. S. and L. Cranor, "Privacy critics: UI components to safeguard users' privacy", *ACM Proceedings of Conference on Human Factors in Computing Systems (CHI'99)*, Pittsburg, Pennsylvania, short paper v.2, pp. 158-159 (1999).

[4]   Adams, A. and M. A. Sasse, "Privacy issues in ubiquitous multimedia environments: Wake sleeping dogs or let them lie?," in *Human computer interaction-INTERACT'99*, A. Sasse and C. Johnson, Eds. Amsterdam: IOS Press, pp. 214-221 (1999).

[5]   Adams, A. and M. A. Sasse, "Users are not the enemy", *Communications of the ACM*, No. 42, pp. 41-46 (1999).

[6]   Adams, A., M. A. Sasse, and P. Lunt, "Making passwords secure and usable", *Springer Proceedings of HCI'97 People & Computers XII*, pp. 1-19 (1997).

[7]   Amalberti, R., "La conduite des systèmes à risques", Presses Universitaires de France, Paris (1996).

[8]   Arief, B. and D. Besnard, "Technical and Human Issues in Computer-Based Systems Security", School of Computing Science, University of Newcastle upon Tyne, UK, Technical Report CS-TR-790 (2003).

[9]   Armstrong, I., "Viruses: Preparing for the Onslaught", *Secure Computing, May issue*, pp. 24-30 (2001).

[10]   Bainbridge, L., "Difficulties in complex dynamic tasks", available online at http://www.bainbrdg.demon.co.uk/Papers/CogDiffErr.html (1998).

[11]   Bastien, C., *Les connaissances de l'enfant a l'adulte : organisation et mise en œuvre*. Paris, Armand Colin (1997).

[12]   Besnard, D. and D. Greathead, "A cognitive approach to safe violations", University of Newcastle, UK, Technical Report CS-TR-791 (2002).

[13]   Bhandari, I. S., H. A. Simon, and D. P. Sieworek, "Optimal probe selection in diagnosis search", *IEEE Transactions on Systems, Man and Cybernetics*, No. 20, pp. 990-999 (1990).

[14]   Briney, A. and F. Prince, "Does size matter? 2002 Survey," in *Information Security*, vol. September 2002 (2002).

[15]   Brostoff, S. and M. A. Sasse, "Are passfaces more usable than passwords?", *People and Computers-XIV Usability or Else! Proceedings of HCI2000*, Sunderland, UK, pp. 405-424 (2000).

[16]   Brostoff, S. and M. A. Sasse, "Safe and sound: a safety-critical approach to security", *Proceedings of New Security Paradigms Workshop*, Cloudcroft, NM, pp. 41-50 (2001).

[17]   Brostoff, S. and M. A. Sasse, "Ten strikes and you're out: Increasing the number of login attempts can improve password usability", *Proceedings of CHI 2003 Workshop on HCI and Security Systems*, Fort Lauderdale, Florida (2003).

[18]   Burkardt, J.-M. and F. Detienne, "La réutilisation en génie logiciel : une définition d'un cadre de recherche en ergonomie cognitive", *Proceedings of ErgoIA 94*, Bayonne, France, pp. 83-95 (1994).

[19]   Cranor, L. F., J. Reagle, and M. S. Ackerman, "Beyond concern: Understanding net users' attitudes about online privacy," in *The internet upheaval: Raising questions, seeking answers in communication policy*, I. Vogelsang and B. Compaine, Eds. Cambridge, MA: MIT Press, pp. 47-70 (2000).

[20]   CSI/FBI, "Computer crime and security survey", Computer Security Institute, Southampton, PA, available online at http://www.gocsi.com/forms/fbi/pdf.html (2003).

[21]   Dekker, S., "Failure to adapt or adaptations that fail: contrasting models on procedures and safety", *Applied Ergonomics*, No. 34, pp. 233-238 (2003).

[22]   Denning, D., *Information Warfare and Security*, Addison Wesley - ACM Press Books (1999).

[23]   Dourish, P., J. D. d. l. Flor, and M. Joseph, "Security as a practical problem : Some preliminary observations of everyday mental models", *Proceedings of CHI 2003 Workshop on HCI and Security Systems*, Fort Lauderdale, Florida (2003).

[24]   DTI-UK, "Information Security Breaches Survey 2002", United Kingdom's Department of Trade and Industry, Technical Report available online at http://www.dti.gov.uk/industry_files/pdf/sbsreport_2002.pdf (2002).

[25]   Evans, S. B. T., C. Harries, I. Dennis, and I. Dean, "General practitioners' tacit and stated policies in the prescription of lipid lowering agents", *British Journal of General Practice*, No. 45, pp. 15-18 (1995).

[26]   Flechais, I. and M. A. Sasse, "Developing secure and usable software", To be presented at OT2003 (March 30th-April 2nd 2003).

[27]   Friedman, B., D. Hurley, D. C. Howe, E. Felten, and H. Nissenbaum, "Users' conceptions of web security: A comparative study", *Proceedings of CHI 2002*, Minneapolis, Minnesota, pp. 746-747 (2002).

[28]   Fujita, Y., "Actualities need to be captured", *Cognition, Technology & Work*, No. 2, pp. 212-214 (2000).

[29]   Furuta, K., K. Sasou, R. Kubota, H. Ujita, Y. Shuto, and E. Yagi, "Human factor analysis of JCO criticality accident", *Cognition, Technology & Work*, No. 2, pp. 182-203 (2000).

[30]   Gasser, L., "The integration of computing and routine work", *ACM Transactions on Office Information Systems*, No. 4, pp. 205-225 (1986).

[31]   Geek.com, "Major Open Source code repository hacked for months, says FSF", available online at http://www.geek.com/news/geeknews/2003Aug/gee20030814021314.htm (2003).

[32]   Gordon, S. and R. Ford, "Cyberterrorism?", *Symantec Security Response* white paper available online at http://www.symantec.com/avcenter/reference/cyberterrorism.pdf .

[33]   Grinter, R. E. and D. K. Smetters, "Three challenges for embedding security into applications", *Proceedings of CHI 2003 Workshop on HCI and Security Systems*, Fort Lauderdale, Florida (2003).

[34]   Hafner, K. and J. Markoff, *Cyberpunk : outlaws and hackers on the computer frontier*. New York, Simon & Schuster (1995).

[35]   Hatch, B., J. Lee, and G. Kurtz, *Hacking Linux Exposed: Linux Security Secrets & Solutions*, Osborne/McGraw-Hill (2001).

[36]   Hoc, J.-M. and R. Amalberti, "Diagnostic et prise de décision dans les situations dynamiques", *Psychologie Française*, No. 39, pp. 177-192 (1994).

[37]   Holmström, U., "User-centered design of secure software", *Proceedings of Human Factors in Telecommunications*, Copenhagen, Denmark (1999).

[38]   Honeynet-Project, *Know Your Enemy: Revealing the Security Tools, Tactics, and Motives of the Blackhat Community*, Addison-Wesley (2001).

[39] Jermyn, I., "The design and analysis of graphical passwords", *Proceedings of the 8th USENIX Security Symposium*, Washington DC (1999).

[40] Leveson, N., "High pressure steam engines and computer software", *IEEE Computer*, No. 10, pp. 65-73 (1994).

[41] Mancini, G., "Commentary: Models of the decision maker in unforeseen accidents", *International Journal of Man-Machine Studies*, No. 27, pp. 631-639 (1987).

[42] McClure, S., J. Scrambray, and G. Kurtz, *Hacking Exposed: Network Security Secrets & Solutions*, Osborne/McGraw-Hill (1999).

[43] Middleton, J., "DoS attack storms port 445", available online at http://www.vnunet.com/News/1131065 (2002).

[44] Mitnick, K. and W. Simon, *The Art of Deception: Controlling the Human Element of Security*, Wiley (2002).

[45] NetworkAssociates, "W32/Deloder.worm", available online at http://vil.nai.com/vil/content/v_100127.htm .

[46] Nielsen, J., "Security and Human Factors", available online at http://www.useit.com/alertbox/20001126.html (2000).

[47] Patrick, A. S., A. C. Long, and S. Flinn, "HCI and security systems", *Proceedings of CHI 2003 Workshop on HCI and Security Systems*, Fort Lauderdale, Florida (2003).

[48] Reason, J., *Human Error*, Cambridge University Press (1990).

[49] Redmill, F., "The COTS question is one of evidence", *Safety Systems*, No. 10, pp. 8-10 (2001).

[50] Redmill, F., "Some Dimensions of Risk Not Often Considered by Engineers", *Journal of System Safety*, No. Q4, pp. 22-40 (2002).

[51] Rescorla, E., "Security holes... Who cares?", *Proceedings of 12th USENIX security symposium*, Washington DC, USA, pp. 75-90 (August 2003).

[52] Richards, D., "The Reuse of Knowledge: a User-Centered Approach", *International Journal of Human-Computer Studies*, No. 52, pp. 553-579 (2000).

[53] Sasse, A., "Computer security: Anatomy of a usability disaster, and a plan for recovery", *Proceedings of CHI 2003 Workshop on HCI and Security Systems*, Fort Lauderdale, Florida (2003).

[54] Sasse, M. A., S. Brostoff, and D. Weirich, "Transforming the weakest link - a human computer interaction approach to usable effective security", *BT Technological Journal*, No. 19, pp. 122-131 (2001).

[55] Simon, H. A., *Models of Man*. New York, Wiley (1957).

[56] Taylor, P., *Hackers: Crime in the Digital Sublime*, Routledge (1999).

[57] Todd, P. A. and I. Bensabat, "The Influence of Decision Aids on Choice Strategies Under Conditions of High Cognitive Load", *IEEE Transactions on Systems, Man and Cybernetics*, No. 24, pp. 537-547 (1994).

[58] Veyrac, H., J.-M. Cellier, and A. Bertrand, "Modèle de l'opérateur et modèle du prescripteur. Le cas des consignes de résolution de situations incidentelles pour les conducteurs de trains", *Le Travail Humain*, No. 60, pp. 387-407 (1997).

[59] Weirich, D. and M. A. Sasse, "Pretty Good Persuasion: A First Step Towards Effective Password Security in the Real World", *Proceedings of New Security Paradigms Workshop*, Cloudcroft, NM, pp. 137-144 (2002).

[60] Whitten, A. and J. D. Tygar, "Usability of security: A case study", School of Computing Science, Carnegie Mellon University, *Technical Report CMU-CS-98-155* (1998).

[61] Whitten, A. and J. D. Tygar, "Why Johnny can't encrypt: A usability evaluation of PGP 5.0", *Proceedings of 9th USENIX security symposium*, Washington DC, USA (1999).

[62] Yee, K.-P., "User interaction design for secure systems", *Proceedings of the 4th International Conference on Information and Communication Security*, Singapore (2002).

[63]   Zetter, K., "Viruses: The Next Generation," in *PC World, online at http://www.pcworld.com/resource/printable/article/0,aid,32802,00.asp* (2000).

[64]   Zurko, M. E. and R. T. Simon, "User-Centred Security", *Proceedings of Workshop on New Security Paradigms*, Lake Arrowhead, CA, pp. 27-33 (1996).