Defining "The Weakest Link": Comparative Security in Complex Systems of Systems

Wolter Pieters
TU Delft and University of Twente
Delft and Enschede, The Netherlands
Email: w.pieters@tudelft.nl

***Abstract*—Cloud architectures are complex socio-technical systems of systems, consisting not only of technological components and their connections, but also of physical premises and employees. When analysing security of such systems and considering countermeasures, the notion of "weakest link" often appears. Humans are then typically said to be the "weakest link" when it comes to security, but no proof is provided for this statement. One reason for this is the fact that there are no unified metrics of security that would apply to physical, digital and social components of complex systems alike. How does one compare the security of a room against the security of a piece of data, and how does social engineering an employee compare to exploiting a server vulnerability? Are we really comparing apples and oranges here, or would it be possible to present a comparative metric that would apply across the different domains? This paper explores the possibility of such a metric for complex systems, and proposes one in terms of the risk induced by an entity in the system. This also provides a foundation for the notion of "weakest link", in terms of the entity (set of entities) with the highest induced risk.**

***Keywords*—Attacker utility, comparative security, induced risk, security metrics, security risk assessment, socio-technical security, weakest link.**

I. Introduction

Information systems as well as cyber attacks become increasingly complex. For cloud architectures, this means that risk management becomes a daunting task. Attacks on information infrastructures may proceed in different stages, where steps occur in the physical or social world, or in a completely different part of the digital infrastructure. Examples include the Diginotar attack, using fake certificates to spy on network traffic [1], and the StuxNet attack, bypassing the air gap with physical transfer of infected USB sticks to sabotage nuclear plants [2]. The complexity of systems and attacks makes it extremely hard to keep track of possible attack paths without appropriate tool support and risk metrics.

Also, human weaknesses may be exploited in attacks, such as when sending phishing mails, or when requesting access to a building dressed like a plumber. In cloud infrastructures, especially multi-step social engineering attacks on credentials may have devastating consequences.[1] It is often said that humans are the "weakest link" in cyber security, and that a system is only as secure as its weakest link. However, the claim that humans are the weakest link has thus far been unsubstantiated, and formal definitions that would enable verification of the claim are lacking.

[1] www.wired.com/gadgetlab/2012/08/apple-amazon-mat-honan-hacking/

In complex systems with complex attacks, support for identifying weak links is essential to assist in security investment. This requires comparison of the security of different entities in a system, physical, digital, as well as human. But for physical, digital, and human security, it is often said that the responsible departments operate as "silos", unaware of the practices on the other side and the consequences for their own tasks. In the world of audits, this may cause security assessments to miss essential interdependencies between the physical, digital, and human aspects of security, and leave the auditor ultimately unable to state where the real weaknesses reside.

This paper addresses the problem how to compare levels of security provided by different components (physical, digital, human) in complex socio-technical systems. We aim at providing a metric for comparing the entities in a socio-technical system in terms of the (in)security they contribute to the system. This will allow us to provide a formal definition for the concept of weakest link. In this paper, we investigate the notion of "weakest link" as a neutral concept, not necessarily referring to humans. Throughout the paper, we use a physical security example for illustration purposes, as physical security is naturally represented on a map (a graph). However, the exact same arguments apply to digital infrastructures and their corresponding graph structures, and modelling formalisms are available for this purpose [3], [4], [5].

In section II, we discuss related work from the information security and risk management communities. In section III, we identify the key components needed for the unified metric. In section IV, we define the metric formally, and in section V we provide examples of using it, in particular within the well-known framework of attack trees. We conclude in section VI, in particular identifying opportunities for future research.

II. Related Work

Several ways to measure security have been proposed in the literature. One way to approach the issue is to classify the risk associated with vulnerabilities. This is for example the case in the Common Vulnerability Scoring System (CVSS) [6], the Common Weakness Enumeration (CWE) / Common Vulnerabilities and Exposures (CVE) databases [7], and the National Vulnerability Database (NVD, http://web.nvd.nist.gov/view/ncp/repository). These systems provide some level of quantification for software security issues, but do not extend to physical or social weaknesses.

It has been proposed to measure security in terms of the "weakest successful adversary", i.e. the weakest adversary that

978-0-7695-5095-4/13 $31.00 © 2013 IEEE
DOI 10.1109/CloudCom.2013.101

would be successful in attacking the system at hand [8]. This is in essence an aggregated difficulty metric over all possible attack paths, specifying the minimal difficulty of getting to the specified goal. However, it does not take the impact for the organisation (or the gain for the attacker) into account, and can therefore not supply a security metric for comparing "weak links".

The weakest link was discussed informally in [9]. The weakest link concept also occurs in game-theoretical approaches to security. In [10], an iterative model of security investment is developed that uses the weakest link concept. However, the focus in all these approaches is on single-asset situations, and therefore the weakest link is simply the attack least costly to the attacker. There is no notion of the distribution and accessibility of multiple assets, for example in case access to the target asset would imply access to another valuable asset. In [11], weakest-link and weakest-target games are used to analyse incentives for security investment. However, the weakest link notion in such analyses abstracts away from the system architecture, and is only expressed in terms of abstract investment.

Considerations similar to the one in this paper have been used in defining the cost-effectiveness of security measures, or return on security investment (see e.g. [12], [13]). In such approaches, the risk in a system is calculated both with and without the countermeasure in place. In this sense, the best countermeasure is the counterpart of the weakest link, but the conceptualisation and formalisation of the latter are new in the present work.

III. Key components

A. Difficulty and risk

As we have discussed above, it is not always easy to align the physical, digital, and social domains of security. In particular, different fields use different variables to express the notion of "security" itself, and this makes it hard to develop models that transcend the silos. Consider the following examples: (1) the security of a door is expressed as the time it takes to open it by force, given certain tools available to a burglar; (2) the security of a cryptographic method is expressed as the key length; (3) the security against social engineering is expressed as the likelihood of an attacker being successful (or: the percentage of people falling for the scam).

In these examples, security is somehow related to the difficulty of performing a step in an attack associated with the security feature at hand. Thus, the notion of difficulty may provide some common ground for unification of metrics. Difficulty can be expressed as a function from attacker investment (time, money, computing power) to probability of success. For the initial presentation of the idea in this paper, we choose to represent it as required attacker investment in terms of money only. For the sake of simplicity, we assume all attacker investment is covered by this value, and we do not consider invested time separately.

Still, difficulty may not cover the intended meaning of weakest link completely. Consider a door protecting € 100 and an identical door protecting € 1000 (Fig. 1). Although the doors provide the same difficulty for the attacker, can we

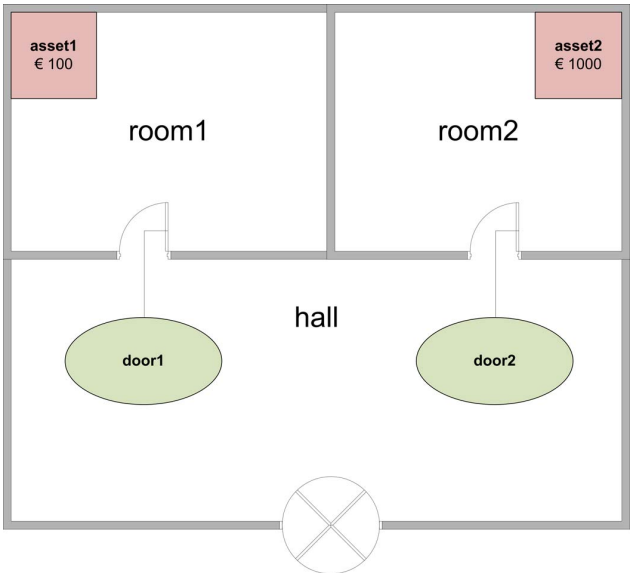

Fig. 1. Two identical doors. Are the doors equally secure, or equally weak links?

say that the doors provide the same level of security, or that the doors are both equally weak links in the system? In other words, does it matter how much an attacker can gain (or the system can lose) by a breach of security? On the one hand, one could say that a security metric should be independent of the impact, as the effort an attacker would have to spend would be the same. On the other hand, a door protecting less valuable assets is less likely to be the target of an attack, and can therefore be considered more secure (assuming the attacker knows what is behind it).

In this paper, we argue for a metric in terms of the risk induced by an entity in the system. This metric takes both the loot and the expected costs (difficulty) of getting there into account (cf. [14]). One can then calculate the risk both with the possibility of getting access through the door and without, and the difference provides a security metric for the door. This approach does not claim anything about the usefulness of the door in a business context. Thus, the door being the weakest link does not necessarily mean that the system should be changed, as legitimate access through the door may be too important.

Risk is typically defined in terms of the likelihood and impact of an event, and we base ourselves on the Risk Taxonomy of The Open Group here [15]. In this paper, we simplify the presentation of the risk concept somewhat. For security risks, the risk concept typically has to take into account an attacker model, because the attacker determines the frequency of threat events, and thereby the risk. Here, this is not our primary concern, and we assume (like in the research on attack trees) that we know the value of the assets for the attacker.[2] We can then approximate the risk by calculating the total expected utility for the attacker when attacking the

[2]Even when there are weak links based on attacker value, the defender may not be interested in protecting assets that have no value for her. We do not explicitly consider value for the defender here.

system. In the above example, if the doors would have no security (and ignoring opportunity costs), the expected utility for the attacker would be € 1100. Security measures on the doors would reduce the utility for the attacker, because he would incur the costs of breaking their security. These costs may involve direct costs such as tools and time, as well as risk of detection and punishment. This is not the place to discuss these in detail, and interested readers are referred to [16]. We do not cover detective controls in the current text.

By quantifying the security in terms of induced risk (or induced attacker utility), this approach also provides the foundations for defining the notion of weakest link. Simply focusing on the effort it takes to get through, like in the above example of a door, does not enable such a comparison. In a socio-technical system, a weak door protecting no valuables cannot be said to be the weakest link, even though it would be easy to get through. Therefore, getting asset value or impact into the equation is essential for defining this concept, and analysing multi-asset situations thus requires different solutions than the single-asset situations discussed in existing literature.

B. Attack paths and system representations

The notion of induced risk of an entity depends on the security of other entities. For example, if an asset is protected by two doors, and it is impossible (or at least very hard) to get through the latter, then the first door does not contribute to the risk, no matter how easy it is to open it (Fig. 2). Therefore, calculating induced risk requires determination of possible attack paths in the system.

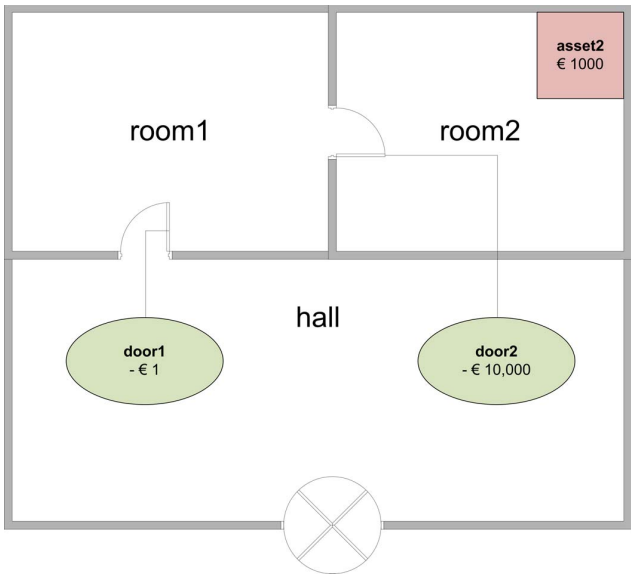

Fig. 2. Two doors in sequence. Because there is no utility for the attacker in trying to break door2, door1 therefore does not induce additional utility, even though its cost of passing is very low.

First of all, a model of the system at hand needs to be available, expressing possible changes in access relations. This is required to distinguish the possible attack paths in the situation where the attacker can get through the door and the situation where the attacker cannot, in order to say something about its security. (This will be discussed in detail further on.) It is also necessary for calculating the expected benefit an attacker can gain. The latter is done by following the possible paths from the attacker position to assets in the system, and accumulating the expected effort needed to reach those. Each step will have an expected effort associated with it, like the time it takes to open a door by force.[3] We assume that the attacker utility associated with effort is known (negative utility), as well as the value of the assets for the attacker (positive utility), and we annotate the models with utility values accordingly.

In this model, we can calculate the expected benefit for an attacker in two cases, when evaluating the security of a certain part of the system, say a door. In the first case, the attacker is allowed to interact with the door. In the second case, he is not allowed to do so. He might try to break in through a window instead, but this may be more difficult. Basically, in the second case, the entity of which we are trying to assess the security is removed from the system, *but rather than implying immediate access, this implies no access at all*. Thus, it is assumed that removal of an entity from the system blocks all attack scenarios that try to get access to something through this entity. This means that no access is the default situation, and it requires explicit entities to enable access. For example, a door is not considered a measure against entering a room, but a means of accessing the room. This is essential for being able to grasp the weakest link concept.

As representations of the system, we use graph-based notations. In particular, we choose the ANKH [5] model for illustration purposes, which consists of entities and groups of entities. In addition to the original qualitative model, we annotate the model with attacker costs on policies. For example, opening the door with a key costs 0, and without a key it costs 800.

C. Attacker models

As mentioned earlier, we do not explicitly consider attacker models in this paper, but assume that the utilities of both breaking security and accessing the assets are known for the attacker under consideration. We also assume full knowledge of these values on the part of the attacker (white-box analysis), and a rational strategy of maximising utility. Attackers will thus adapt their strategies based on the knowledge they have about the system. In the case of a fully rational and omniscient attacker, the door only contributes to risk if the *easiest* path to an asset goes through the door, and this attack provides positive utility to the attacker. However, in possible extensions involving probabilistic attacker models, where attackers might not always take the "best" path, the measures would become more complicated.

In the version of the metric proposed here, we use the so-called monotonicity assumption [17]: going back to a place already visited costs 0 for the attacker. This makes the analysis easier, as one does not have to consider cycles in the state space. Lifting this assumption could be an extension for future work, but we assume the metric would not become much more

[3] In more advanced models, this will typically be a probability distribution, where the expected effort is the mean effort needed to get through.

accurate. Firstly, costs of going back are typically low, and secondly, costs of going back would have to be taken into account for any step, both in the case with and without the entity present in the system, thereby evening out any deviations in the metric. Only in cases in which the attacker has a choice between a step where it is very hard to go back in real life and one where this is very easy may such extensions prove useful.

Another question is how to handle situations in which there are multiple ways to pass a certain security feature. For example, an attacker who fails to open a door with a screwdriver may attempt different approaches. When there are multiple ways in, should the expected utility increase accordingly? If one way fails, the attacker can try another (if not caught). However, in most situations, the different ways in require other actions first, such as getting the key or social engineering someone. Therefore, this is not a local concern of the door only. The metrics assume that *any* interaction with the door is disabled. Therefore, the metrics will take into account both trying to break the door and using a previously acquired key. The metrics thus consider how much risk the door contributes to the system as a whole, independent of whether other parts of the system make it easier to get through the door (e.g. by acquiring a key through social engineering).

IV. Defining the metrics

In this section, we provide definitions that enable calculations of the weakest link in a complex system. The central concepts are assets, guardians, utility and induced risk.

Definition 1: An *asset* is an entity in the system that has positive utility to the attacker when gaining access. The positive utility is called *reward*.

Definition 2: A *guardian* is an entity in the system that the attacker could try to pass to gain access to an asset. The *cost* of passing a guardian determines the negative utility for the attacker when deciding to pass. The cost is typically dependent on the entities an attacker already has access to, such as keys or passwords

We can make use of a type / group system (or generally: an instance / subclass relation), such that the metrics can be applied to individual entities as well as all entities of a specific class. For example, one could say that Bob has a certain level of induced risk, but also that all humans together have a certain level of induced risk. Of course, when speaking about the weakest link, one needs to define clearly what is being compared to what. It does not make sense to say that all humans have a higher induced risk than each individual computer or each individual door, and that therefore humans are the weakest link.

Definition 3: A *guardian group* is a set of guardians, typically defined based on some common property.

Definition 4: An *attack suite* is a partially ordered set of actions to pass guardians. From an attack suite, the set of compromised assets can be calculated.

Definition 5: The *utility* associated with an attack suite is calculated as the maximum of (a) 0, and (b) the rewards of the compromised assets, minus the costs of passing the guardians. The *maximal utility* of an attacker with respect to the system is the maximal utility that an attacker can gain by passing guardians and gaining access to assets. The maximal utility can be calculated as the maximum over the utilities of all possible attack suites.

Definition 6: The *induced risk* associated with a guardian or guardian group is the maximal utility in the system including the guardian (group), minus the maximal utility in the system with the guardian (group) removed. The *weakest link* in a system is the entity that has the highest induced risk.

If the induced risk is negative, there is no incentive for an attacker to try to get access via that entity. The higher the induced risk, the more reason for an attacker to take the step of gaining access via the entity. When multiple similar guardians exist, like in many people being targeted in a phishing attack, one may want to consider them as a guardian group for the purpose of calculating the induced risk. This is what we mean when we say that "humans are the weakest link": it is not necessarily the case that one human, when taken individually, is weak, but rather that, when targeting a lot of humans, the overall expected utility is higher than in the case of targeting technology.

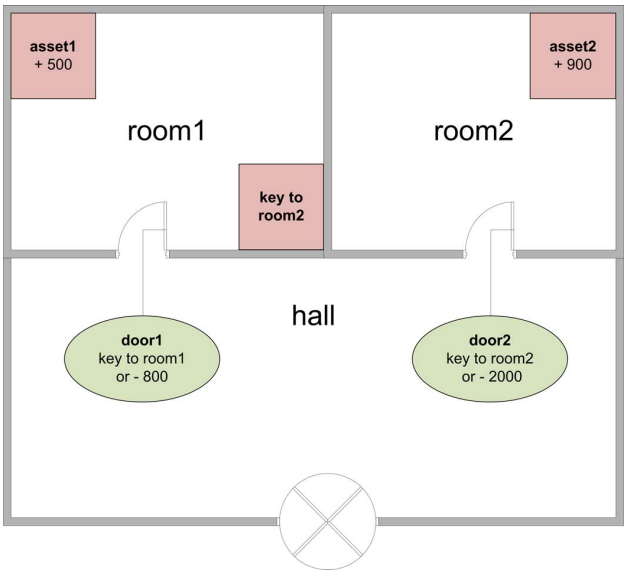

Fig. 3. An example system map with two doors and two assets.

In Figure 3, a simple system map is shown with two doors and two assets. Again, for easy illustration, we use physical access mechanisms in this example, but the metrics work for digital examples (e.g. passwords and web services) as well as for keys and doors. For `room1`, an attacker needs either the key to get in, or it will cost him 800 value units. For `room2`, an attacker needs either the key, or incurs a cost of 2000 value units. The asset in `room1` is worth 500 units, and the asset in `room2` 900. A key to `room2` is present in `room1` at no cost. There is no key to `room1` accessible to the attacker. We now calculate the expected utility for the attacker in the following situations: (1) the system as shown; (2) the system without the door to `room1`; (3) the system without the door to `room2`.

In the first case, the best option for the attacker is to spend 800 to get access to `room1`, get the asset and the key, open

room2, and get the asset from there as well. The utility for the attacker is -800 + 500 + 900 = 600. In the second case, there is no way to get into room1, and therefore there is no way to get the key to room2 either. The only thing the attacker can do is open room2 by force, incurring a cost of 2000, and getting an asset worth 900. As this path yields negative utility, the attacker will choose not to do this, and therefore have utility 0 overall. In the third case, room2 is not accessible, and, for similar reasons, the attacker will decide not to open room1, again yielding utility 0. This means that both doors have an induced risk of 600 in this system.

If the asset in room1 were worth 900 instead, the expected utilities for the attacker would be 1000, 0, and 100, respectively. In this situation, the induced risk of the door to room1 is 1000, and that of the door to room2 900. In this case, the door to room1 is the weakest link. Note that the fact that there is a key in room1 is of central importance here, illustrating the relevance of dependencies in the system. When we add the human element, we could for example include help of Alice as a credential for access to room1. The attacker could then invest in acquiring the assistance of Alice in opening the room for him. The difficulty of this step is obviously essential for determining whether Alice is the weakest link in this system or not. Within the TRESPASS project[4], we investigate the inclusion of social science results in the risk models, linking behavioural research on compliance with risk management concepts.

Until now, we have calculated the induced risk for single entities. However, the claim "humans are the weakest link" applies to a *group* of entities. The definition is analogous here: remove all entities in the group from the system, recalculate the expected utility for the attacker, and determine the difference. This is *not* necessarily the same as summing up the values for the individual entities in the group, as dependencies may exist between them.

The metrics can be implemented as part of security (investment) models, as discussed in the next section.

V. Using the metrics

A. Attack trees

Application of the metrics is foreseen in existing formalisms and tools within information risk management. One of the most prominent formalisms for security analysis involves calculations on attack trees [16], [18], [19]. Basic attack trees only represent attacker actions, and do not associate those with the parts of the system that are the target of the attack step. For example, an attack tree would have a node "acquire password by phishing", but this does not explicitly state that humans are the target of this step, i.e. humans will give the attacker access to the password. In order to use the weakest link metrics in an attack tree context, we therefore need to annotate actions with the entities (guardians) that are the target of the actions. In principle, this is only needed for the leaf nodes. Guardian groups can be used as labels as well as guardians (see Fig. 4).

When each action (leaf node) in an attack tree has been annotated with an entity, as well as the values of interest, one can start the calculation of the metrics. One would first calculate the metrics for the original attack tree. For each entity, one then constructs an attack tree in which all nodes associated with the entity are removed. In this new attack tree, one calculates the same risk metrics, and the difference is the risk induced by the entity. In Fig. 4, the attacker utility for the original tree would be 1000 - 500 - 100 = 400, when taking the cheapest path. Now consider the situation where all social nodes are removed. As the left part of the tree is an AND-node, this whole branch would become impossible. Instead, the attacker would be forced to take the right branch. The utility would now be 1000 - 700 = 300, meaning that the induced utility of the social guardian group is 100. The same holds for the physical guardian group, as removal blocks the left branch of the tree as well. As removal of the digital group does not affect the optimal attack path, the induced utility for this group is 0.

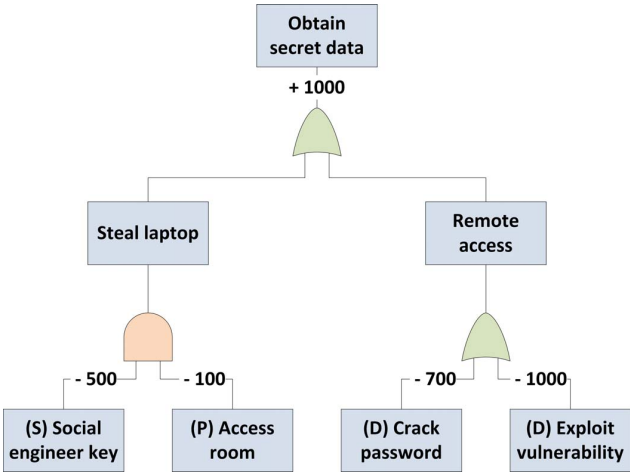

Fig. 4. An example attack tree. Physical nodes are labelled (P), social nodes (S), and digital nodes (D). The labels thus represent guardian groups.

As many quantitative analysis techniques and tools for attack trees already exist, they form an easy platform for implementing the metrics, in combination with annotations of the attack trees with the corresponding guardians or guardian groups. However, attack trees are still limited to single-asset situations, and do not cover dependencies between steps.

B. Other opportunities

a) Attack navigators: The need for annotating attack trees with the corresponding model entities is another argument for not considering attack trees as a separate formalism, but rather to something that is related to a "map" of an organisation, in which the actions of the attacker take place [5]. Ideally, the metrics could be applied directly on the system maps rather than on attack trees generated from those. In that case, situations where multiple assets can be targeted appear naturally in the model, whereas in principle one attack tree would be needed for each asset. In the TRESPASS project, we are currently investigating this direction using the concept of attack navigator. In this approach, we also aim at including explicit attacker models rather than having implicit information on the attacker in the model of the system, such as for example skill level. The cost of passing a guardian would then be

[4] www.trespass-project.eu

dependent on both the resistance provided by the guardian (control strength) and the skill and tools of the attacker (threat capability). This approach is inspired by the Risk Taxonomy of The Open Group [15].

b) Insider threat: In insider threats, the attacker will already have initial access to part of the system, for example a person working as a technician within a cloud provider. In this case, the starting point for the analysis is different, and different entities may turn out to be weak links. This is because the outside perimeter may be a strong link (low induced risk), but in case of insider threat that does not help much. For insider threat, the method would thus be able to identify weak links *inside* the organisation, as the assumption is that the attacker already made it through the outside perimeter(s).

c) Penetration testing: One of the hard questions in risk assessment is how to validate the results of proposed methods. Real attacks may not happen often enough to provide significant validation, and controlled testing may therefore be needed. Estimations of weak links may be verified in penetration tests, and may at the same time provide guidance for where to direct the tests. There are challenges involved in socio-technical penetration testing [20], but they can be essential in verifying the estimations used in the models. They can also be used to update the difficulty estimations.

VI. Conclusions

This paper provides, to the best of our knowledge, the first formal definition for what constitutes the weakest link in a socio-technical system, with respect to security risk. It does so by providing a security metric in terms of the risk induced by an entity in the system, basically comparing the risk in the system with the entity present and without it. The attack paths on which the metric is based guarantee inclusion of dependencies in the system, and also make it possible to apply the metric directly to attack trees. Moreover, the metric takes into account multi-asset situations, where the weakest link may not simply be the easiest or cheapest path to a specified asset.

The metric provides an important step towards guiding security investment in complex systems of systems, such as cloud architectures. When it is known which type of component, location or person contributes most to the security risk, one can indeed direct investments towards the weakest link, which is not possible based on an informal concept only. Cloud infrastructures will be one of the case studies for applying the metrics within the attack navigator concept in the TRE$_S$PASS project. In future work, we aim at integrating the approach outlined here with explicit attacker models, such that different metrics can be calculated depending on assumptions on the goals, skills, and resources of attackers interested in the system. In particular, one could introduce attackers that do not always choose the cheapest path, giving rise to more complex, but potentially also more inclusive, risk metrics.

Acknowledgement

The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement ICT-318003 (TRE$_S$-PASS). This publication reflects only the authors' views and the Union is not liable for any use that may be made of the information contained herein.

References

[1] N. Leavitt, "Internet security under attack: The undermining of digital certificates," *Computer*, vol. 44, no. 12, pp. 17 –20, 2011.

[2] R. Langner, "Stuxnet: Dissecting a cyberwarfare weapon," *Security & Privacy, IEEE*, vol. 9, no. 3, pp. 49 –51, 2011.

[3] T. Dimkov, W. Pieters, and P. Hartel, "Portunes: representing attack scenarios spanning through the physical, digital and social domain," in *Proceedings of the Joint Workshop on Automated Reasoning for Security Protocol Analysis and Issues in the Theory of Security (ARSPA-WITS'10)*, ser. LNCS, vol. 6186. Springer, 2010, pp. 112–129. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-16074-5_9

[4] C. W. Probst and R. R. Hansen, "An extensible analysable system model," *Information security technical report*, vol. 13, no. 4, pp. 235–246, 2008.

[5] W. Pieters, "Representing humans in system security models: An actor-network approach," *J. of Wireless Mobile Networks, Ubiquitous Computing, & Dependable Applications*, vol. 2, no. 1, pp. 75–92, 2011.

[6] M. Schiffman, G. Eschelbeck, D. Ahmad, A. Wright, and S. Romanosky, "CVSS: A common vulnerability scoring system," 2004.

[7] R. Martin, "Common Weakness Enumeration (CWE v1.8)," *National Cyber Security Division, US Dept. of Homeland Security*, 2010.

[8] J. Pamula, S. Jajodia, P. Ammann, and V. Swarup, "A weakest-adversary security metric for network configuration security analysis," in *Proceedings of the 2nd ACM workshop on Quality of protection*, ser. QoP '06. New York, NY, USA: ACM, 2006, pp. 31–38. [Online]. Available: http://doi.acm.org/10.1145/1179494.1179502

[9] I. Arce, "The weakest link revisited," *Security & Privacy, IEEE*, vol. 1, no. 2, pp. 72–76, 2003.

[10] R. Böhme and T. Moore, "The iterated weakest link," *Security Privacy, IEEE*, vol. 8, no. 1, pp. 53–55, 2010.

[11] J. Grossklags, N. Christin, and J. Chuang, "Secure or insure?: a game-theoretic analysis of information security games," in *Proceedings of the 17th international conference on World Wide Web*, ser. WWW '08. New York, NY, USA: ACM, 2008, pp. 209–218. [Online]. Available: http://doi.acm.org/10.1145/1367497.1367526

[12] B. Blakley, E. McDermott, and D. Geer, "Information security is information risk management," in *Proceedings of the 2001 workshop on New security paradigms*. New York, NY, USA: ACM, 2001, pp. 97–104. [Online]. Available: http://doi.acm.org/10.1145/508171.508187

[13] H. Cavusoglu, B. Mishra, and S. Raghunathan, "A model for evaluating it security investments," *Commun. ACM*, vol. 47, no. 7, pp. 87–92, 2004. [Online]. Available: http://doi.acm.org/10.1145/1005817.1005828

[14] S. Hakim, G. F. Rengert, and Y. Shachmurove, "Target search of burglars: A revised economic model," *Papers in Regional Science*, vol. 80, pp. 121–137, 2001. [Online]. Available: http://dx.doi.org/10.1007/PL00013617

[15] The Open Group, "Risk taxonomy," The Open Group, Tech. Rep. C081, 2009. [Online]. Available: www.opengroup.org/pubs/catalog/c081.htm

[16] A. Buldas, P. Laud, J. Priisalu, M. Saarepera, and J. Willemson, "Rational choice of security measures via multi-parameter attack trees," in *Critical Information Infrastructures Security*, ser. LNCS. Springer, 2006, vol. 4347, pp. 235–248. [Online]. Available: http://dx.doi.org/10.1007/11962977_19

[17] P. Ammann, D. Wijesekera, and S. Kaushik, "Scalable, graph-based network vulnerability analysis," in *Proceedings of the 9th ACM Conference on Computer and Communications Security*. New York: ACM, 2002, pp. 217–224. [Online]. Available: http://portal.acm.org/citation.cfm?doid=586110.586140

[18] S. Mauw and M. Oostdijk, "Foundations of attack trees," in *Proc. 8th Annual International Conference on Information Security and Cryptology, ICISC'05*, ser. LNCS, D. Won and S. Kim, Eds., vol. 3935. Springer, 2006, pp. 186–198. [Online]. Available: http://www.icisc.org/

[19] B. Schneier, "Attack trees: Modeling security threats," *Dr. Dobb's journal*, vol. 24, no. 12, pp. 21–29, 1999.

[20] T. Dimkov, A. van Cleeff, W. Pieters, and P. Hartel, "Two methodologies for physical penetration testing using social engineering," in *Proceedings of the 26th Annual Computer Security Applications Conference*. New York, NY, USA: ACM, 2010, pp. 399–408. [Online]. Available: http://doi.acm.org/10.1145/1920261.1920319