An Asset to Security Modeling? Analyzing Stakeholder Collaborations Instead of Threats to Assets

Andreas Poller
Fraunhofer SIT
Darmstadt, Germany
poller@sit.fraunhofer.de

Sven Türpe
Fraunhofer SIT
Darmstadt, Germany
tuerpe@sit.fraunhofer.de

Katharina Kinder-Kurlanda
GESIS – Leibniz-Institute for the Social Sciences, Köln, Germany
katharina.kinder-kurlanda@gesis.org

ABSTRACT

Risk assessment in information security traditionally analyzes threats to assets. An asset is a persistent item or property of value and has an owner. Attacks damage assets; security controls prevent attacks to preserve their value. Expected attack loss is calculated from the value of the attacked assets. This common analytic approach works satisfyingly if an IT system runs in an enclosed environment within an organization. Nowadays, IT systems are accessed and used across organizational boundaries by a multitude of independent stakeholders employing them for their own interests and with particular expectations regarding their trustworthiness. The asset paradigm cannot support estimating consequences of security incidents that may harm these complex stakeholder collaborations. We propose to model the stakeholder collaboration networks and to analyze scenarios of how security incidents affect relationships between stakeholders. Collaboration continuously creates value for all participants. Security incidents change the behavior of stakeholders and their willingness to collaborate, but in complicated ways. Transmission factors characterizing a relationship help us to assess the impact of incidents. We apply the conventional method and our new approach to a case study and compare the results.

Categories and Subject Descriptors

H.1 [**Information Systems**]: Models and Principles; H.1.2 [**Models and Principles**]: User/Machine Systems—*Human factors*; K.6.5 [**Management of Computing and Information Systems**]: Security and Protection

General Terms

Security, Human Factors

NSPW'14, September 15–18, 2014, Victoria, BC, Canada.

http://dx.doi.org/10.1145/2683467.2683474.

Keywords

Assets, collaboration networks, harm analysis, requirements elicitation, risk assessment, security engineering, threat modeling

1. INTRODUCTION

Assets are a key concept within information security – they are what we protect. To understand the impact of an attack, however, we need to look beyond assets. The literature conceptualizes assets broadly as "something someone places value upon". Asset identification and asset relationship analysis play important roles as starting points for risk analysis because they allow estimating potential losses in the case of security incidents. We argue that, despite the broad definition in place, assets are really considered in a much narrower way, namely as sensitive entities or artifacts that need protection, have owners and carry value by themselves. This perspective reflects the roots of information security and its original purpose to protect military, government or business organizations from the adverse outside world. Assets are internal entities – information, systems, or physical items – that should remain internal.

However, contemporary IT systems operate in the context of complex multi-stakeholder collaborations. Organizations are no longer closed entities but rather comprise actors engaged in processes of cooperation within complex networks. System architectures are equally complex and often distributed amongst several organizations. To estimate security risks to a system operator or software provider we nowadays need to consider relations and tensions between the various system stakeholders. For many IT security aspects, we can no longer draw a line between the "restricted outside" and the "permitted inside" where mere border crossing causes harm to assets. Instead of considering a well-defined organization with few overall interests we are now dealing with various stakeholders seizing and incorporating IT systems for their purposes – all with different backgrounds, with different interests that need to be considered in the IT systems' design, and with complex relations of power and influence. Following Freeman we also see stakeholders as anyone, be it group or individual, "who can affect or is affected by the achievement of the organization's objectives" [8, p. 46].

Stakeholder interests and the resulting complexity in relations of power and influence have already been addressed by scholars by including new, less tangible assets such as *reputation* or *trust* – often called abstract assets. We argue

that these attempts are merely patches to a concept fundamentally flawed. The asset concept per se is misleading if we are to achieve useful descriptions of the security challenges of modern IT systems which are deployed into multilateral usage scenarios. Instead of adapting the concept of assets to attempt a description of such multilateral interests, we propose to identify and analyze the stakeholder collaboration networks that make use of IT systems. In our security analysis of these networks we dismiss the concept of assignable values because in this context values are not firmly attached to an entity or artifact. Rather, stakeholders continuously *generate value* for each other through their interactions, in particular by using the IT system. Hence, instead of trying to assign values to assets, we propose to develop and analyze potential impact scenarios in which stakeholder cooperation might degrade, and to identify factors contributing to or preventing such degradation.

We suggest that by exploring stakeholder collaborations we can better identify and monitor the security risks to a system and can more easily propose adequate risk mitigations to the system owner. As we see the concept of assets and values and our concept of collaboration networks and degradation scenarios as incommensurable, our intent is not to abandon the asset concept but to restrict its use to where it fits best: as a means to estimate risks to nameable entities or (information) artifacts that have value primarily to internal stakeholders within an organization.

In this paper we contrast the traditional asset paradigm in risk assessment with our new stakeholder collaboration paradigm. Traditional risk assessment as described in ISO 31000 goes through several stages of analysis to identify appropriate risk treatments. Our approach assumes that risk treatments are determined by collaboration relationships, and that security requirements follow from expected reactions to incidents. We illustrate the differences between the traditional and our paradigm using the example scenario of the security of a research data archive.

First we describe the asset concept as commonly understood and apply it to analyze risks in our scenario. Then we introduce the stakeholder collaboration paradigm and outline an approach to relationship-centered risk assessment. Revisiting our scenario, we look at its risks from this new angle. We compare the results of both approaches and show how the stakeholder collaboration paradigm yields insight that cannot be attained in the conventional approach. Finally, we show that our approach can also lead to a better understanding of security risks in other real-world scenarios by using the example of stakeholder collaborations around a hotel reservation portal.

2. SOCIAL APPROACHES TO SECURITY

Dourish and Anderson [6], considering security technologies as means to ensure privacy, argue for a social reading of security. They observe that although the literature acknowledges the social origin of privacy and security concepts, most treatments of privacy take a narrow view based on an economic rationality model. Dourish and Anderson propose to see security and privacy as both practical action and as discursive practice in order to capture their role in the management of social relationships and not just in the management of information. Our paper explores the implications of such a view for secure systems engineering.

2.1 Social Production of Security and Privacy

Recent work on peer-produced security and privacy [4, 9] addresses the challenge of producing security or privacy as a common resource in communities. Economic theory has established conditions for effective governance of common goods, without which individually rational behavior would have detrimental, sub-optimal social outcomes. Ostrom's framework [18, 5] posits five such conditions: (1) resources and their use can be monitored; (2) rates of change are moderate; (3) there are dense social networks and frequent communication; (4) certain parties can be excluded from using a resource; and (5) resource users support monitoring and enforcement. Using these criteria one can analyze and design institutional arrangements to increase the chances of effective governance [9]. Social context and social processes can thus be a source of oversight and security. A classic example is separation of duty [22] in such a way that multiple parties must collaborate to complete a task. Lipford and Zurko [15] generalize this idea and propose to integrate community oversight into security mechanisms. Multilateral security [19] emphasizes another social aspect of security, namely balancing conflicting needs in the design of security mechanisms.

We take a different position in this paper: We start from a given organizational arrangement and analyze the consequences of security incidents. We do not assume stakeholders to be *peers*; to the contrary, we analyze their respective reactions to security incidents and the underlying factors.

2.2 Security Economics

Real-world relationships can also have an adverse impact on security. Anderson [1] notes, for example, that security technology can be used to shift blame and liability from one party to another, rather than just to make everyone more secure. However, we do not yet fully understand the impact of institutional arrangements and social relationships on the economics of security. Cyber-insurance is an example where relationships matter. Böhme and Schwartz [3] have developed a framework for models of cyber-insurance comprising five key components: network environment, demand side, supply side, information structure, and organizational environment. They aim to understand why cyber-insurance has not taken off in the market despite promising theoretical work.

Common to their framework and our approach is the basic assumption that agents extract value from the network. However, our work highlights the complexities of the network environment that arise when we consider application services and the stakeholder networks involved in their operation. Our qualitative treatment suggests that network topologies are complicated and non-standard and that defense functions pertain not to agents but rather to particular relationships between them. Since we identify agents with network nodes, the line between the network environment and the organizational environment becomes blurred. For example, in our approach, an insurer would be modeled as just one node in the stakeholder network. Ours is not an alternative model but rather we take a different viewpoint. Insurance is a means for large communities to distribute and thereby reduce individual risk. While some stakeholder networks may be large enough for insurance to be considered as a means of organizing them, our case study example operates on a smaller scale.

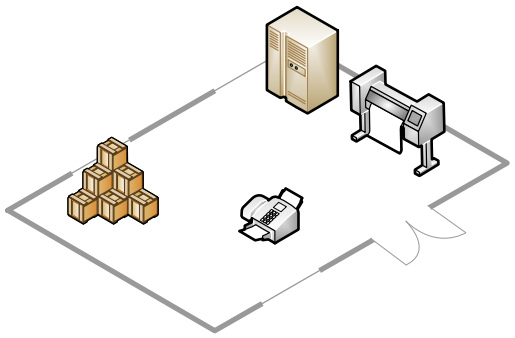

Figure 1: Physical office assets (in the style of [23, Figure 3-3]).

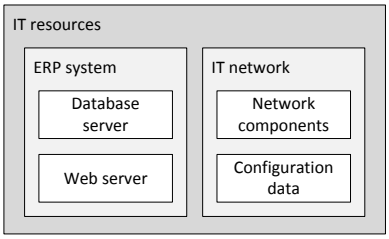

Figure 2: Part-of asset relations. Inner boxes represent components.

2.3 Security Versus Collaboration

Often reality turns out to be more complicated than security policy models assume. Security also causes frictions, especially when it gets in the way of collaboration. The CSCW community picked up these frictions even before the discipline of usable security emerged [25]. Two sources of friction receive particular attention. First, collaboration makes it necessary to reconcile the security requirements of multiple parties [19]. Second, often collaboration and security are competing goals [24]. Collaboration requires easy sharing of information, while security aims to prevent access unless authorized. The search for the sweet spot of easy and secure sharing goes on [12]. Our proposal approaches this challenge from a different angle. We argue that collaboration not only imposes additional requirements on security policies and mechanisms to make systems work, but rather is in itself something to be protected.

3. THE ASSET PARADIGM

In a nutshell, an asset is anything that has value and thus requires protection [7, 10, 11]. The literature seems to agree upon this very broad definition. However, the definition does not reflect how we perceive the asset concept's use in information security practice. In the following, we will explain why widely accepted relations of the asset concept to other concepts narrow down the actual meaning of the term and limit its usage.

3.1 Is It an Item, a Property, or Something Else?

What exactly is an asset: a physical item, an information item, or a property of such items? The Common Criteria provide examples for all of these categories: "contents of a file", "authenticity of votes cast in an election", or "access to a classified facility". Swiderski and Snyder [23, Ch. 3] specifically highlight "physical assets" such as business equipment, as assets that can be relevant for a security analysis (see Figure 1).

All these references to assets are convincing, because they point to something – such as a data record or the confidentiality of particular information – that *can be protected by applying security controls.* While they are different, all these asset types have in common that we perceive their properties as persistent and stable if shielded from external influence. Metaphorically speaking, we can put an asset in a drawer for some time and when we take it out again we assume that, while the asset may have aged, its essential properties are still unchanged – they can be preserved.

The perception of assets as having preservable properties relates to the fact that asset stakeholders have specific *expectations* of these properties and aim to actively control them for their purposes: an IT system operator does not want a file content changing in an undesired or unexpected way; a government board wants to count only those votes in an electronic election that have been cast by legitimate voters. Security controls reflect these interests and the notion of assets helps security engineers to express and model the purposes of security controls.

3.2 Asset Relationship Modeling

In the course of a security analysis it is not only relevant which assets can be affected but also how existing asset relationships may cause harm to "spread" to other assets. Without a model of how harm can propagate it is difficult to estimate all consequences of a security incident and to understand the possible impact on stakeholders' interests.

IT security practice knows several ways of analyzing asset relations. A common approach for creating an asset-relation model takes into account the *structure* of an IT system, in some cases as a part of its environment or application context. To develop *part-whole asset models*, one assumes that harming a sub-component, considered a sub-asset, affects those assets assigned to related superior components. Each sub-component or super-component should be perceived as a single asset in itself as we can reasonably assume that it may have properties worthy of protection. However, some assets also relate to others because of the architecture of a system. Figure 2 illustrates this view.

Anderson's account of multilateral or compartmented security [2, Ch. 9] illustrates this train of thought. Sensitive information in centralized systems constitutes a more valuable asset than the same information distributed over many independent locations; at the same time, more users get access to the system. Compartmentation aims to reduce the exposure of assets by isolating users from each other, making available to each user only the information and resources required for a particular role or task.

Another possibility for modeling asset relationships is using *taxonomies* to manage different levels of abstraction at which assets could be defined, or inheritance relationships among assets. In such models, assets describing superior concepts, e.g., "credit card information", can be refined by defining different child assets, e.g., "embossed digits and let-

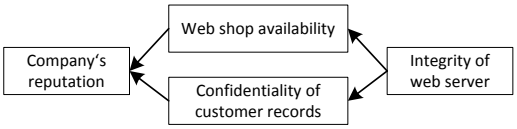

Figure 3: Example: informal asset relations model. Arrows show possible harm propagation.

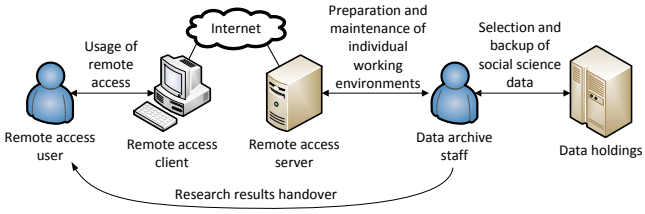

Figure 4: The planned remote access system.

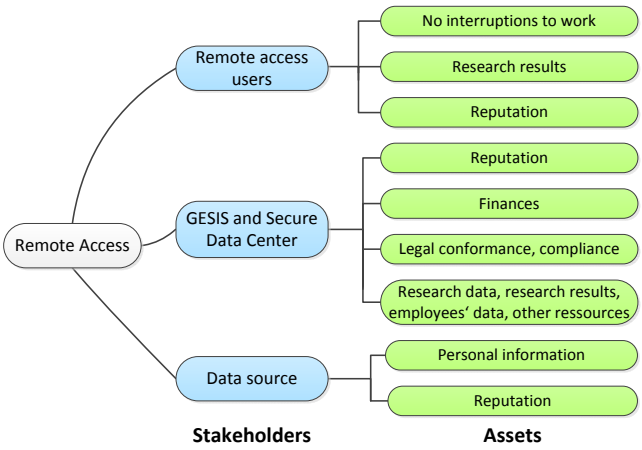

Figure 5: Stakeholders-assets mind map (example).

ters on a plastic card", "data stored in the memory of an on-card chip", and "server's memory data content".

Less formal models capture only expected harm relations between assets without considering in detail why these relations exist. Figure 3 shows an example. Because they are less formal, these models are appropriate for quickly analyzing and depicting different stakeholder perspectives. In Figure 3, "integrity of web server" could hint at a technician's perspective in a company, whereas the management level is worried about the "company's reputation".

In our survey of asset models we found that usually both *tangible* and *intangible* assets were being considered. Tangible assets are those where the perception of the asset, what it consists of, and what its properties are, is to some degree consistent among stakeholders. The most tangible assets are usually physical items, data assets, or properties of items or data. Intangible or abstract assets are concepts such as privacy or reputation. Their meaning often varies and also depends on the individual context of the model's user.

4. SECURITY ANALYSIS WITH ASSETS

4.1 Case Study: Secure Data Center

To illustrate our argument we draw on one particular example of security engineering. Our case study concerns setting up a remote access system at a social science data archive in Germany. The GESIS data archive provides long-term preservation for survey and other research data. Curated, processed and documented data sets are made available to the scientific community for secondary analyses [21]. Researchers can in this way use the data generated by another research project to produce new research results. Some data sets require privacy protection due to their potential for subject re-identification. These data sets cannot be offered for download to the general scientific community.

The GESIS Secure Data Center allows authorized researchers to work with sensitive data. It ensures data protection through various technical, organizational, and contractual safeguards. The current Secure Data Center implementation requires that researchers work in a safe room on the premises of GESIS. A remote access facility is under development. Before getting access to any data, researchers must prove a legitimate research interest and sign a contract specifying the acceptable use policy. The Secure Data Center provides each user with an isolated, individually tailored work environment, which contains only the data sets released to this user and the necessary programs to process the data. Work results can be saved inside the personal environment, from where GESIS Secure Data Center staff may release analysis results to the user after verifying compliance with the policy and contract [14].

While the Secure Data Center does not solve the problem of re-identification itself, it restricts data access to a well-vetted community of researchers, reduces the exposure of sensitive data through compartmentation, and controls data processing and data release. The safe room with its physical access control contributes to this security design. Offering remote access as shown in Figure 4 poses the challenge of maintaining the same level of security while opening an attack surface to the Internet. This was the starting point of our project.

4.2 What Asset Identification Contributes

Conventionally one identifies, analyzes and evaluates risks in order to decide in the end which risk treatments are appropriate. In our first attempt to analyze the security risks for the remote access we followed this approach and tried to identify assets and asset stakeholders. To collect the necessary information, we involved the GESIS development team comprising both application experts familiar with social science data archiving and provision, and IT experts familiar with implementing IT solutions. We asked team members to individually gather (a) stakeholders they considered as having an interest in the security of the remote access and (b) assets that these stakeholders were likely to have an interest in. We requested participants to collect and visualize their thoughts in mind maps; Figure 5 shows an example from one participant. We aggregated all mind maps and identified five stakeholder groups for the remote access IT system:

- **Remote access users**, who perform secondary analyses of the provided data,

Stakeholders	Assets
Remote access users	Privacy, intellectual property, ability to fulfill contract with Secure Data Center, reputation, resources, own research data, analysis results
Data archive with Secure Data Center	Remote access service, reputation, IT infrastructure (material, supply, manpower, time), being compliant with privacy best practice, ability to fulfill contracts with primary investigators, archive holdings
Survey participants	Privacy
Primary investigators	Reputation
Research community	Reputation, participants' willingness to enter surveys

Table 1: Collected stakeholders and assets.

- **Data archive**, which the Secure Data Center is a service of,
- **Survey participants**, whose responses are provided by the data center,
- **Primary investigators**, who conduct surveys and provide data to the archive, and
- **Research community**, as a conceptual and regulatory umbrella above all activities.

Our participants collected a multitude of assets, see Table 1. During a later group discussion it became clear that all assets were relevant and justifiable from the individual perspectives of team members.

In order to understand which security measures would be reasonable for the remote access we performed a risk analysis. We collected – in addition to identifying assets and stakeholders, and again in collaboration with the development team – a set of four primary threats pertinent to the Secure Data Center. Primary threats were defined as potential adverse actions that could harm the interests of at least one of the stakeholders. The collected threats were:

- **Violation of confidentiality** of social science data by remote access users or external attackers, including re-identification of an individual;
- **Theft of intellectual property** by remote access users or external attackers obtaining research data, results, records or whole databases;
- **Denial of service or sabotage** towards the Secure Data Center systems by external attackers;
- **Preparation of other attacks** by external attackers, using the Secure Data Center systems as a host.

The list could be questioned, modified, extended and refined. It will however suffice as our working threat model to illustrate the points we wish to make in the following. Therefore we would like to save the discussion of proper threat modeling techniques for another paper.

In the next step we developed scenarios of how each threat could harm assets from our list of threats. The CORAS [16] model in Figure 6 illustrates the process: a remote access user may attempt to de-anonymize study data provided to her in the course of research activities at the Secure Data Center. She may succeed and identify an individual participant of the initial survey, thus gaining access to personal information. This event would harm the asset *confidentiality of personal information*. According to development team

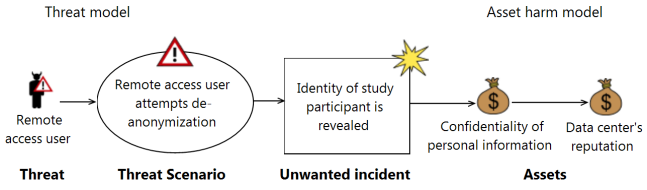

Figure 6: How a threat might affect assets related to the remote access in our example scenario.

members this would have negative effects on the *data center's reputation*.

The example in Figure 6 shows how threats and assets relate to each other and the role they play in security analyses: The diagram shows a cause-effect relationship rooted at the threat agent. In the diagram we show a particular way in which the threat agent might attack the remote access; we model *how* the threat could manifest. The chain of effects following the threat agent's actions leads to an unwanted incident. Actions we perceive as a security breach lead to a state undesired by the system's stakeholders.

The unwanted incident in Figure 6 is a demarcation line: the left side of the model shows *who* could attack the remote access, and *how* this attack could take place. To the right of the incident details are shown of the problem we face if the incident occurs. Arrows pointing towards the incident denote actors and actions to reach the state, arrows pointing away from the incident describe negative effects once the state has been reached. Traditional risk analysis uses such a view to estimate the risk that the assumed threats pose to the operation of the IT system. To this end, analysts usually combine both sides: the model of actions leading to the incident serves the purpose of estimating *how likely* the unwanted states occur. The model of the incident effects is used to estimate the *extent of the damage* once the unwanted state has been reached.

Estimating asset risks promises to allow weighing one risk against another and defining priorities for finding, evaluating and selecting security controls such that the overall risk remains within acceptable limits. In practice risk analyses are often messy because lack of knowledge, uncertainty and hidden assumptions can influence analytic results [13]. However, that is not the issue we want to focus on. Our discussion is concerned with the analytic means for investigating consequences of security incidents and thus involves the issue of what asset concepts can contribute to such an analysis.

4.3 Where Asset-based Risk Analysis Fails

At first glance, risk analysis boils down to an analytic combination of information about *how likely* a security incident is to occur and *what harm* it could do. Consequently, risk is often expressed with the equation:

$$Risk = Likelihood\ x\ Impact$$

Let us assume that we can somehow estimate the likelihood of events, e.g., by using statistical data on how often certain attacks usually take place and what may motivate attackers to specifically target the remote access. Estimating the potential consequences of a security incident through asset analysis is more difficult since it is hard to oversee

values and related costs in case of harm for some of them. Consider the following example: if a security incident causes harm to the Secure Data Center's asset *IT infrastructure* we can develop scenarios *how* the IT infrastructure could break down under an attack, and can then establish *who* would be affected by this breakdown. Given this information, we can calculate how many working hours may be lost due to reduced productivity of affected employees, and how many working hours it may require to switch over to redundant systems, fix affected network components, re-establish processes, and so forth.

Risks to the asset *IT infrastructure* can be analyzed in the conventional way: a security incident affects an asset under control of the organization running the IT system. This organization has specific policies governing who has access to the asset, under which circumstances, and for what purpose. An attack on the IT system violates this policy, and the asset as a resource for the organization's businesses loses desired properties. The loss can be estimated considering what it takes to re-establish desired properties and to compensate irreversible damages.

Yet, the analytic starting point is strikingly different in two situations we encountered in our case study. First, the situation is different in those cases where the Secure Data Center and an external player create an asset collaboratively. This is the case for assets such as *reputation* and *trust*. An organization cannot "produce" or "buy" these intangible assets to exploit them as resources. The organization can act or cooperate in a particular way in the hope that others attribute trustworthiness or a good reputation to it. Security incidents may affect this attribution process. The resulting damage to the organization cannot be estimated with the information we collected so far. Second, the situation is also different if an organization does not own an asset. This is the case for assets like the *intellectual property*, which belongs to the remote access user. Even if we could calculate the loss to the remote access user if this asset is harmed, the user's loss does not necessarily translate into loss for the system operator, namely the Secure Data Center.

4.4 Could Trust Modeling Help?

Our first attempt to analyze the remote access security risks consisted of identifying assets that may be affected and may take damage by attacks on the IT system. We isolated those assets to which we can assign expected costs should they be damaged because they are owned by the Secure Data Center, and which are kinds of physical or information artifacts. For other assets we cannot do so because they are owned by a stakeholder other than the Secure Data Center, are too abstract, or both.

What ensues if a security incident affects these assets? Two main consequences are to be expected: first, once a security incident affects interests of an external stakeholder, it will influence this stakeholder's future expectations of the system and its provider. The system and its operator become less trustworthy due to not operating in the way the stakeholder expects; this deviation interferes with the stakeholder's interests.[1] Second, if an external asset is harmed, the stakeholder may require direct compensation from the system operator or software vendor. However, this consequence is rare in practice because it is hard to address and enforce, e.g., through lawsuits. Thus, we focus primarily on the first issue, the impact on trust relationships.

Figure 7 shows trust relationships between the remote access's stakeholders. One significant trust chain begins with the survey participants and ends at the remote access users. Each party trusts the next in this chain to handle personal information appropriately. Trust relations between the Secure Data Center and the remote access users are mutual: remote access users also trust that the Secure Data Center properly handles their research results. The trust model in Figure 7 shows how trust relationships could degrade due to an attack. For example, if attackers violated the intellectual property rights of a remote access user by getting unauthorized access to research results (cf. the threats discussed in Section 4.2), the user would probably question the Secure Data Center's trustworthiness.

But does modeling of trust relationships help to estimate the security risks for the Secure Data Center? We argue that this is in fact not the case unless we learn more about what could actually happen if stakeholders' expectations are not met. To this end, we have to take one further step: we need to analyze how the stakeholders collaborate.

5. COLLABORATION RISK ANALYSIS

5.1 Effects on Stakeholder Collaborations

Models of trust relationships allow localizing where tensions between trustors and trustees may arise once expectations are not met, but they do not enable describing potential consequences of trust breaches. However, we need exactly the latter for our risk analysis. So let us assume that security incidents occur and stakeholders reconsider their perception of other stakeholders' trustworthiness. What could be the result? If we take into consideration that the trustor's leverage is primarily (and often only) the existing collaboration with the trustee, it is reasonable to search for possible consequences there. In general we can identify four impact scenarios:

1. **Lethargy:** The collaboration is *not affected* and no consequences ensue because the trustor remains unaware or has no other options.
2. **Withdrawal:** The trustor *reduces the level of collaboration* with the trustee over time, perhaps after condoning a security incident at first glance. A subtler form of withdrawal is subverting the trustee and undermining her interests.
3. **Intervention:** The trustor tries to reestablish trust by *intervening in the trustee's business* with the aim to reinforce trustworthy practices there.
4. **Compensation:** The trustor requires compensation in addition to the other options. As already mentioned, this card is rarely drawn.

Withdrawal, intervention or compensation may go against the interests of the system operator or software provider. They can damage her business and may have undesired monetary consequences. However, harming external stakeholders' assets does not necessarily imply these impact scenarios. While a security incident, or the perceived risk of a security incident, might trigger an impact scenario, interdependencies and circumstances related to a stakeholder collaboration work as counterforces and might prevent that an impact scenario manifests.

[1] If a stakeholder expects an insecure system, the incident is no problem. Expected insecurity is already priced into the collaboration.

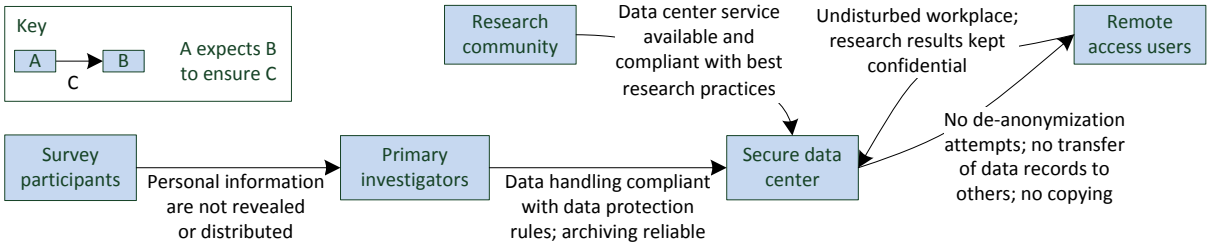

Figure 7: A model of trust relations between stakeholders of the remote access.

The four impact scenarios correspond loosely with the four common risk responses: acceptance, avoidance, mitigation, and sharing or transfer [17]. The difference is that we consider impact scenarios to be a consequence of the properties of a relationship. While decisions about risk responses are the result of a conventional risk assessment process, impact scenarios are to be analyzed as part of the process and security requirements follow from the expectation of undesired responses.

5.2 Reaction Effort and Disruption Costs

The notion that harmful effects of security incidents may materialize differently depending on the circumstances and scenarios is already important for asset-centric risk analyses. While a security incident can cause harm to assets, there may be technical or organizational means or circumstances limiting its negative effects. An example of such a means is compartmentalization by a network firewall separating a computer network into an intranet and a demilitarized zone. The firewall can prevent that compromising a web server affects employees' internal workstations. If the firewall separation is missing, damage could spread much more easily and might put more assets at risk.

For stakeholder collaborations, the different *scopes of actions* of the stakeholders involved in a collaboration play a crucial role in how harm spreads as a consequence of a security incident. The likelihood that one of the four impact scenarios becomes manifest is primarily influenced by the stakeholders' individual abilities and interests to influence and shape how a collaboration takes place. We see stakeholders' scopes of actions as limited by two main factors: *reaction effort* and *disruption cost*.

Reaction effort refers to the fact that withdrawing from using a software, intervening with the software provider or vendor, or asking for compensation requires effort of a trustor, i.e., collaboration partner. The trustor has to spend time and money to enforce one or multiple scenarios, for example, if she intervenes to enforce demands or claims.

Disruption costs are caused by the *reciprocity* of the collaboration between the software provider and the affected collaboration partner. Disruption costs become effective if the collaboration partner decides to withdraw from a service: by lowering the activity level in a collaboration the collaboration partner also loses the benefits from the collaboration, just as the trustee does.

Possible reaction effort and disruption costs require tradeoffs from external stakeholders: addressing harm caused by a security incident may represent additional disadvantages that outweigh the perceived damage. This tension influences decisions of those affected by security incidents with regards to how they exploit their position to forward damage to a software vendor or service provider, or what actions they take to prevent incidents in the future.

Without a thorough analysis of the stakeholders' scope of actions and the reciprocity between stakeholders and collaborators, the risk analysis will deliver misleading answers: we may tend to overestimate consequences of security incidents in the case that we underestimate the conditions restraining collaborators from addressing harm. Or we may underestimate possible consequences in the case that we lose track of collaborators' means and engagement to interfere with the system operator or software provider.

5.3 Modeling Interdependencies

Modeling interdependencies between stakeholders is necessary to reflect on how violating collaborators' interests can rebound on the one responsible for security. Such reflection requires information about whether and how going against collaborators' interests may or may not provoke the four impact scenarios mentioned above. We call these pieces of information the *transmission factors*. These factors could be properties of the software or system, as well as conditions or circumstances underlying the collaboration of the IT system operator or provider with an external stakeholder, or means collaborators have to enforce their interests. Typical examples of transmission factors are:

- **Legal situation:** Are claims or demands against the software vendor or provider enforceable?
- **Ambiguous culpability:** Was the security incident partially caused by actions of stakeholders other than the responsible operator or provider, e.g., the software's users?
- **Competitive situation:** What are the stakeholder's alternatives to using the software affected by the security incident? Are there competitors who offer a more secure software with similar or even better functional properties? Has the software provider a monopoly?
- **Lock-in:** What are the costs for two parties to suspend the collaboration and to move to other collaboration partners with similar services?
- **Collaborator interconnectedness:** How strongly is an affected collaboration partner connected to other actual or potential collaborators and can she influence their decisions to the disadvantage of the software provider or operator?
- **Liability towards third parties:** Does the collaboration partner have liabilities towards others that may be affected by a security incident?
- **Risk tolerance:** To what extend do stakeholders accept imposing additional risks on their business?

- **Cost-benefit equilibrium:** Are the costs and benefits of using a software or a service in a precarious equilibrium so that security incidents easily lead to stakeholders perceiving costs as considerably outweighing benefits. Maintaining the cost-benefit equilibrium is often an issue for high-margin commercial services or software offerings.

Explicating transmission factors has a further benefit besides supporting reasoning about the likelihood and extent of potential security incident consequences: once identified, transmission factors can be monitored to detect risk-relevant changes in the conditions underlying the system's operation or the provision of software. System operators or software vendors can react early by adapting their software and processes.

5.4 Refining Stakeholder Identification

In the course of analyzing stakeholder collaborations and transmission factors influencing how security incidents lead to impact scenarios, the set of identified stakeholders may require further refinement because of two reasons: first, the analysis might reveal that a middleman (e.g. consultants, lawyers, authorities, other organizations) between a stakeholder and the IT system operator actually constitutes a unique stakeholder in its own whose interests, expectations and scope of actions need to be considered and treated separately. In that case a middleman should appear as a separate stakeholder in the collaboration network model. To provide an example for the Secure Data Center, a data protection commissioner could be conceived as a middleman whom inquiring institutions might employ to enforce data protection requirements towards the data archive. However, depending on the scenario, data protection commissioners could also be unique stakeholders if they initiate actions towards the data archive on their own.

Second, we may discover that a more fine-grained analysis is required for stakeholders that aggregate and thus abstract a group of individual stakeholders with similar interests and involvement in the collaboration network. For instance, a stakeholder "customer" often aggregates a multitude of individual customers for (necessary) model abstraction. But in the course of a transmission factors analysis, we might come to the conclusion that, e.g., customers' vendor dependency differs substantially among this group, and so does the likelihood that their individual collaborations with the system operator or software vendor degrade after security incidents occur. In this case, we need to further unpack aggregated stakeholders during the transmission factors analysis, for example, by defining different classes of sub-stakeholders according to the varying manifestation of one or more transmission factors.

The challenge of finding the right level of abstraction for modeling stakeholders applies to asset-centric analyses too. In practice, it is often useful to start with a very small set of more generic and aggregated stakeholders and to refine them iteratively if it turns out that this could affect the modeling outcome.

5.5 Proposed approach

We have explained how analyzing intangible assets or assets of external stakeholders contributes little to risk analysis. Although security incidents may affect intangible assets, effects do not directly translate into consequences for responsible system owners, operators or software vendors. Instead security incidents may affect collaboration links to the external stakeholders while the strength of the impact depends on individual transmission factors. Out of this understanding of security risks in multi-stakeholder collaborations we propose a risk analysis approach different from those approaches relying on asset identification. Our approach comprises six immediate steps plus one advanced step:

1. **Model the stakeholder collaboration network** for the IT system considering the software or system provider and external stakeholders that trust in the security of the IT system or software. Note the stakeholders' expectations of the security of the software or IT system. Also model collaborations between external stakeholders.
2. If a threat model is available, **assign threats and resulting potential security incidents** to those stakeholders whose interests they may go against. Mark those collaborations that may be affected by this particular interest being harmed.
3. **Collect transmission factors** and their characteristics for those collaborations that may be negatively affected by potential security incidents. Consider that stakeholders may employ other stakeholders to address their harm towards the system or software provider.
4. Use transmission factors to **estimate** for each collaboration **how likely** it is that **a security incident** as a manifestation of a threat **will lead to each of the four** general **impact scenarios**. Also search for indications that an impact scenario can be ruled out.
5. **Prioritize impact scenarios according to their relevance** to the stakeholder collaboration network and the possible harm to the interests of the system or software provider. To this end, analyze the downsides that degrading collaborations might have for the system or software provider and her business and combine it with the likelihood that a security incident provokes the impact scenarios. Also consider assets of the system provider not covered by the collaboration analysis.
6. **Use impact scenario relevance for risk assessment and development decisions**. Depending on the current status of development, information about impact scenario relevance can be used to perform a further risk assessment or to render some first decisions in an early stage of development. To perform a risk assessment a further analysis of potential attack paths is necessary which requires that at least a first prototype of the software or IT system is available.
7. In the aftermath: **monitor transmission factors** to see whether conditions change and security models need to be reconsidered.

All suggested steps can be executed in the course of a group process of definition involving application and domain experts within an organization. It can be useful to repeat early steps if subsequent steps reveal new insights not considered before.

6. EXEMPLARY WALKTHROUGH

Going back to the Secure Data Center scenario we show in the following how our approach can be used in practice. Our description covers all steps but, as it is intended as an exemplary walkthrough only, misses some analytical details

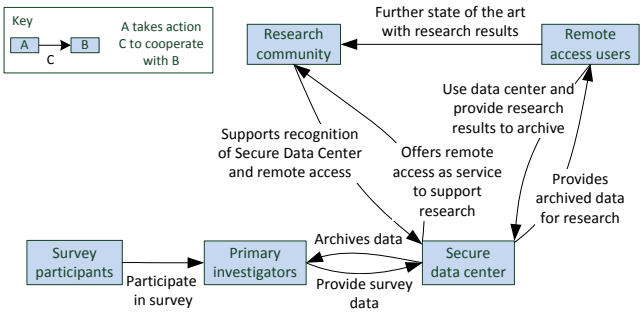

Figure 8: Secure Data Center remote access collaboration network.

Stakeholders	Threats				
	Personal information disclosure	Intellectual property theft	Denial of service/ sabotage	Prepare attacks	
Survey participants	•				
Primary investigators	•	•			
Research community	•	•	•		
Remote access users		•	•		
Secure Data Center			•	•	

Table 2: Threats and affected stakeholders.

that might be relevant to fully understand the actual security risks for the Secure Data Center.

6.1 Modeling the Collaboration Network

The first step aims at understanding the collaboration scenario in the intended IT system setting, including stakeholders directly using or having access to the IT system, and stakeholders otherwise relevant in the overall ecosystem. In many cases this step might not be much different from identifying asset stakeholders. For the Secure Data Center scenario we can re-use the stakeholders we have already collected in Table 1.

If we followed a conventional asset analysis we would now start collecting the assets belonging to the stakeholders. However, our analytic approach suggests describing instead how the stakeholders cooperate. Such cooperation often comprises reciprocal actions beneficial to all involved parties. For example, the Secure Data Center provides data to the remote access users enabling them to conduct their data analyses. On the other hand, by using the remote access, users contribute to the reputation and recognition of the Secure Data Center.

By collecting stakeholders and collaborations we can depict the stakeholder collaboration network. This can be done using a graphical notation as is shown for the Secure Data Center collaboration network in Figure 8. Because of the depicted dependencies of the stakeholders, the network provides indications where interruptions in the network could negatively impact our application scenario. For example, if survey participants hesitate to cooperate with primary investigators, the data archive, which the Secure Data Center is a part of, receives less data or data of less good quality. Stakeholder expectations are not shown in Figure 8; they were already provided in Figure 7.

6.2 Assigning Threats

In a conventional asset analysis we would now analyze how threats may affect the collected assets. This is usually done by investigating the potential asset harm caused by materialized threats, i.e., unwanted states or security incidents an attacker might provoke.

However, we modeled the setting for our security analysis not by means of assets and their owners but with stakeholders and the collaborations between them. As in the conventional analysis of how threat manifestations may harm assets, we first investigate how security incidents might infringe on stakeholders' interests, and then how the infringement can affect the collaborations because of stakeholders' reactions. We call these effects on stakeholder collaborations impact scenarios and they are what we need to consider.

For example, in the Secure Data Center a theft of intellectual property might result in a situation where survey data from primary investigators has been transmitted to unauthorized persons. Such a security incident would infringe (1) on the interests of the survey participants because they consider the incident as a risk to their privacy, (2) on the interests of the primary investigators because it could reduce participants' willingness to participate in future studies, and (3) on the interests of the research community because the incident might put into question research and data dissemination practices. Overall there are three potential stakeholders affected and their collaborations with the Secure Data Center put at risk. The other three threats we identified in our first threat analysis affect different stakeholder interests, as shown in Table 2. Note that the threat of preparation of further attacks does not affect an external stakeholder but solely the Secure Data Center and its staff.

After we have identified which threat manifestations, i.e., potential security incidents, affect which stakeholder interests, we can attend to the effects to collaborations caused by stakeholders reacting to the incident and its undesired results. At this point in the analysis it is not important which of the impact scenarios lethargy, withdrawal, intervention, or compensation are more or less likely. As a first step, we strive to generally understand where threat manifestations pose risks to collaborations.

For example, if personal information is leaked by the Secure Data Center, the research community could reduce their support for the remote access. Figure 9 shows an analysis of how other threats may affect collaborations in the Secure Data Center scenario.

6.3 Analyzing Transmission Factors

By means of the previous steps we have produced a theory of how threats may affect stakeholder interests and may consequently affect stakeholder collaborations. We have not yet considered how likely the impact scenarios are, but only whether they are generally possible. To decide on security measures for the Secure Data Center, however, we need to somehow weigh the risks of threats.

In an asset analysis we would elaborate on the value of assets and on the likelihood that an asset is affected temporarily or permanently. In the stakeholder collaboration

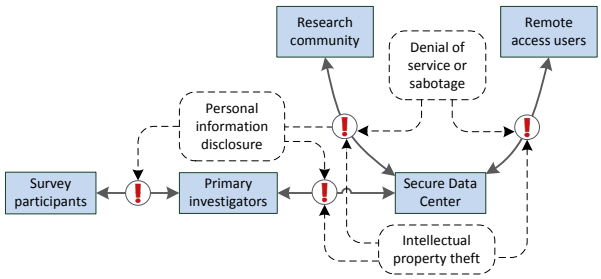

Figure 9: Threats and endangered collaborations.

	Threats								
	Personal information disclosure			Intellectual property theft			Denial of service/ sabotage		
Impact scenarios	W	I	C	W	I	C	W	I	C
Collaboration with									
Study participants	+	-	-	-	-	-	-	-	-
Primary investigators	++	+	+	+	+	+	-	-	-
Research community	+	++	-	+	+	-	+	+	-
Remote access users	+	-	-	+	-	+	++	-	+

Relevance: -/green—low, +/yellow—medium, ++/red—high
Impact scenarios: W-Withdrawal, I-Intervention, C-Compensation

Table 3: Relevance of impact scenarios (lethargy scenario excluded as no treatment is required).

analysis we already know that a threat might lead to one of our four impact scenarios (lethargy, withdrawal, intervention, or compensation) but we need further information to estimate how likely it is that they manifest. To this end, we collect transmission factors, i.e., factors that contribute to or interfere with the manifestation of an impact scenario. Transmission factors can be considered as arguments to justify why we accept an impact scenario as likely or unlikely.

In our example case we elicited transmission factors in collaboration with the application experts of the Secure Data Center. For example, we discussed whether a leak of personal information as a security incident might lead to a withdrawal of primary investigators from a cooperation with the Secure Data Center. It turned out that some primary investigators are obliged by their research grants to provide their research data to an archive or similar institution. Consequently, withdrawing is not very likely for these research partners. It also turned out that there is an institutional interweaving of important players in the research community with the institution's advisory boards or with other powerful bodies. Hence, the application experts were concerned that primary investigators or other collaborators in the research community could exploit these relationships to intervene in the Secure Data Center. In this way we identified two transmission factors: *obligations to cooperate with the data center* and *institutional interweaving*.

6.4 Prioritization of Impact Scenarios

Transmission factors enable us to reason about how likely it is that a threat (having become manifest as an attack) will provoke one of the four impact scenarios. Finally, to get an estimate of how relevant the impact scenarios are, we need to investigate how severely an impact on a collaboration will affect the interests of the service or software provider, in our case the Secure Data Center's interests. To this end, we need to develop hypotheses about the extent to which collaborations might degrade and how degradation may disrupt the business of the IT system operator.

For example, in our discussion with the application experts, it became apparent that interventions by the research community could have a serious impact on the Secure Data Center's operation because of its funding scheme. Similarly, if primary investigators would withdraw from the service, the Secure Data Center remote access would lose relevance leading to limited research funding opportunities.

To estimate the relevance of an impact scenario out of this information we need to consolidate possible impacts with the likelihood that impact scenarios affect collaborations. As described in the previous step the impact scenarios of withdrawing primary investigators and intervening research community players can likely lead to a degradation of collaborations with the Secure Data Center in the case that private information of participants is leaked via the remote access.

Mapping the conclusions about likelihood and possible impact to an overall relevance of an impact scenario can be done in several ways. Each conceivable mapping may have a different focus and viewpoint regarding the security problem. For the Secure Data Center we have chosen (a) to generally assign a low relevance to unlikely scenarios, no matter what impact the scenario might have, (b) a medium relevance to scenarios we perceive as both possible and as having moderate, but still manageable consequences, and (c) a high relevance to those scenarios we perceive as both possible and as potentially having a severe, medium, or long-term impact on the Secure Data Center's business. Given this mapping, we can, e.g., conclude that the leakage of personal data of study participants is a scenario of high relevance for the Secure Data Center remote access and needs further consideration. The relevance ratings of impact scenarios for the other threats is shown in Table 3.

6.5 Contribution to the Risk Assessment

We have developed some theories about which threats to the remote access can provoke impact scenarios that might lead to harm to the collaborations between the software or system operator and the other stakeholders. We can use these theories for a first ranking of impact scenarios and related threats according to their relevance.

For a final risk assessment, we would now need to investigate the ways in which a threat could actually lead to a security incident as a pre-stage of an impact scenario. To this end, we would have to turn our attention to the possible technical design, architecture and implementation of the remote access and to the embedding organizational processes, and would have to define the potential attack paths available for adversaries. For example, with respect to the threat of leaking study participants' personal information, one attack path might be attempting to readout data of the remote access user's client systems. Other attack paths could involve compromising communication channels between client and

server at the Secure Data Center, or breaking into network services.

Given these attack paths we could further refine the estimates of the likelihood that attacks take place and that impact scenarios become manifest. Finally, we could calculate the risks to the remote access given a particular technical and organization solution. For our scenario a final risk assessment is not possible because the remote access is still in an early development stage. But the relevance of impact scenarios and related threats we estimated in the previous step helped to focus development efforts and to decide on the design of the system. In addition, our analysis could also be a starting point for developing a stakeholder management strategy [20] for the Secure Data Center.

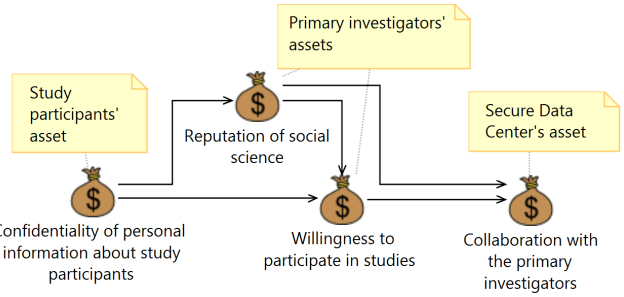

Figure 10: Attempted asset relationship model.

7. COLLABORATIONS VERSUS ASSETS

One could argue that the analytic process and the concepts we propose are merely complex refinements of the asset concept, asset identification and asset relations. Indeed a mapping is possible between some parts of the asset approach and the stakeholder collaboration approach. However, with regards to the goal to anticipate potential consequences if threats manifest, the stakeholder collaboration network analysis takes a unique perspective hard to achieve with the traditional analytic techniques. In the following we will elaborate similarities and differences of the two concepts in more detail revisiting our case study for illustration.

7.1 Capturing Collaborations With Assets

In the collaboration network model we assume that stakeholders have specific interests and expectations of the collaboration with the IT system operator or software provider; the potential results of malicious attacks oppose these interests and expectations. If we wanted to capture this information in an asset model, we would start with an initial asset that a security incident could directly harm, for example the confidentiality of personal information as shown in Figure 6. This initial harm does not have to belong to a stakeholder collaborating with the IT system operator or software provider, but harming the identified initial asset may affect collaborations between these parties. We could define the potentially affected collaboration as an asset in itself – an asset that may be harmed within the course of damages to other assets.

Figure 10 shows an example for the Secure Data Center scenario. Attacks against the remote access may infringe on the privacy of study participants. This privacy breach may reduce the willingness to participate in studies and could affect the reputation of the social sciences overall, and eventually may lower the activity level for the cooperation. We could also add some further assets to this model, such as "freedom from collaborator's interventions" or "freedom from collaborator's claims" to represent the impact scenarios *intervention* and *compensation*.

7.2 Limits of Asset Models

So what is lacking in an asset model as shown it in Figure 10? We propose that the problem lies in the model's "middle section": the asset *collaboration* may suffer harm if the willingness of potential participants decreases or if the reputation of the social sciences is somehow affected. Both asset harms are sufficient conditions for negative effects on the collaboration. However, it is also a reasonable assumption that the primary investigator may already question the collaboration in case there is a relevant risk that the "willingness to participate" could be negatively affected by current security incidents, or even regarding the expectation of further security incidents. How can we express this possibility with assets and asset relations?

By using the asset concept we can model causalities like "harm to asset A may cause harm to asset B" after identifying assets and analyzing their relationships. We assume that modeled causalities are defined through a particular organizational or technical structure of the IT system and its environment. But the situation is strikingly different if there is an independent party with own goals, expectations and risk perceptions "mediating" between an initial effect (such as a privacy breach), and possible consequences (such as a lowered collaboration level).

This party can make individual choices about when and how to react and is exposed to tensions constraining her scope for reactions. Modeling intangible assets such as *reputation*, *trust*, *customer confidence*, and so forth, merely expresses the difficulty of appropriately considering these factors when using current security modeling techniques. Intangible assets and their relations deliver little information as they are mostly generic, and potential consequences in case of harm are held up to arbitrary interpretations.

Collaboration network analysis complements conventional threat modeling with an assessment of harm mechanisms. While asset and threat models help us understand possible attacks, the analysis of collaborations and stakeholders gives us a clear picture of the consequences a security breach might have. We can thus more adequately assess risks and decide on where mitigation efforts are most needed.

Our experiences working with the data center application experts highlight another advantage. Conventional threat modeling requires technical knowledge and security expertise, which the participants of a requirements elicitation process often lack. Collaboration analysis provides application domain experts with a more familiar and intuitive framework to discuss security requirements. While it is hard even for trained professionals to develop a sound asset-stakeholder model, most people are able to explain who they collaborate with, how their relationships work, or how others and they themselves would react to a certain type of incident. Indeed we witnessed how the data center staff, when asked to list assets, instead expressed their concerns with regards to the collaborations with external stakeholders vital to their service domain.

Nonetheless, asset modeling remains useful in cases where

we comprehensively understand, or even have the possibility to influence, how effects on one asset could cause effects elsewhere. This is often the case when the system operator also owns the considered assets. For example, regarding a possible sabotage attack on the Secure Data Center's IT we can in detail investigate the possible damage inside the organization, and can influence how harm could be prevented from spreading. But outside this application domain, the analytic approach of asset identification and asset relationship modeling leads to very generic models lacking information to actually understand the security problem we face.

8. A FURTHER EXAMPLE

In addition to the Secure Data Center example we outline in the following a second real-world scenario where applying our analytic approach can lead to a better understanding of security risks.

The example concerns a hotel reservation portal. Typically such portals involve a multitude of stakeholders such as portal users, hotels, the portal operator, and payment processors, e.g., credit card companies. Stakeholder collaborations are manifold: portal users browse the collected information on hotels, room rates, and customer reviews. The portal operator benefits from customer bookings and gains a small commission. Hotels provide information and receive room bookings. Payment processors provide payment options and receive a commision for each transaction. From the portal operator's viewpoint one important question is what security measures are required to protect from undesired consequences of security threats.

Defining assets worthy of protection such as customers' personal information, credit card information, portal software, server, and hotel information, is only of limited use to the portal operator because it is difficult to value and prioritize these assets. In terms of general value, credit card information might be most valuable, perhaps followed by the IT systems. Processed customers' personal information and hotel information are difficult to value because it is not quite clear what would happen if an attack compromised their confidentiality, availability or integrity. From the asset perspective, for example, the hotel ratings stored in the portal's database might be a valuable asset required by customers for selecting and booking hotels. When attacking the portal an adversary might feed the portal with multiple negative hotel ratings at a large scale. The consequences on the collaboration between portal operator and hotel depend on how responsibilities are shared between both. If hotels are responsible for checking their ratings for anomalies it is unlikely that they blame the portal operator. Other transmission factors are also relevant: do other portals provide a better protection against faked ratings, e.g., by asking reviewers for a proof of having visited the hotel, or by using improved CAPTCHA mechanisms to lock out automated bots? Questions concerning the portal operator's liability may also play a role.

Regarding the impact of potential attacks on credit card information the question for the portal operator is again not primarily one of their value because she does not own this information asset. The value of credit card information might attract adversaries, but the operator may not be affected by a data breach in proportion to their value. The more important question is that of the effects on collaborations. Whether customers will withdraw from the portal if their credit card information is leaked and misused depends on the credit card company's compensation promises, and whether customers can actually connect a misuse with the use of the credit card at the portal. Regarding the cooperation between the payment processor and the credit card company, both might also have no particular interest in abandoning their cooperation with the portal because they receive a commission for every payment transaction, and the popularity of the payment scheme depends on the amount of available acceptance points. The payment provider in particular depends on a large network of collaborators. A more common scenario is intervention: credit card companies nowadays require payment processors to fulfill certain standards to participate in the payment system, for example, the Payment Card Industry Data Security Standard (PCI DSS).

9. CONCLUSION

We have proposed a stakeholder collaboration paradigm for modeling security requirements of IT systems. We have contrasted this new concept with the traditional concept of asset modeling and illustrated the differences using the example scenario of the security of a research data center for the social sciences. While asset modeling is useful for defining threats to internal assets that have singular and internal ownership, it does not easily allow to assess threats posed to stakeholder relationships and externally owned assets. Stakeholder collaboration modeling does allow for this and in addition allows to gauge and classify possible consequences of attacks. We have shown how our paradigm yields insight that either cannot or are very difficult to be attained using the traditional asset approach.

Acknowledgments

We thank our shepherd Daniela Oliveira and all other NSPW participants for the valuable discussion that led to many fruitful contributions to this paper. We also thank Philipp Holzinger, Stefan Triller, and Laura Kocksch for their terrific work in our research project. This research was funded by the European Center for Security and Privacy by Design (EC-SPRIDE).

10. REFERENCES

[1] R. Anderson. Liability and computer security: Nine principles. In D. Gollmann, editor, *Computer Security — ESORICS 94*, volume 875 of *LNCS*, pages 231–245. Springer Berlin / Heidelberg, 1994.

[2] R. J. Anderson. *Security Engineering: A Guide to Building Dependable Distributed Systems*. Wiley, 2nd edition, 2010.

[3] R. Böhme and G. Schwartz. Modeling cyber-insurance: Towards a unifying framework. In *Proc. WEIS*, 2010.

[4] L. J. Camp. Reconceptualizing the role of security user. *Daedalus*, 140(4):93–107, Fall 2011.

[5] T. Dietz, E. Ostrom, and P. C. Stern. The struggle to govern the commons. *Science*, 302(5652):1907–1912, 2003.

[6] P. Dourish and K. Anderson. Collective information practice: Exploring privacy and security as social and cultural phenomena. *Human-Computer Interaction*, 21:319–342, Sept. 2006.

[7] B. Fabian, S. Gürses, M. Heisel, T. Santen, and H. Schmidt. A comparison of security requirements engineering methods. *Requirements Engineering*, 15(1):7–40, 2010.

[8] R. E. Freeman. *Strategic Management: A Stakeholder Approach*. Pitman, 1984.

[9] V. Garg, S. Patil, A. Kapadia, and L. J. Camp. Peer-produced privacy protection. In *Proc. IEEE Symposium on Technology and Society*, ISTAS. IEEE, 2013.

[10] D. Gollmann. *Computer Security*. Wiley, 3rd edition, 2011.

[11] C. B. Haley, R. Laney, J. Moffett, and B. Nuseibeh. Security requirements engineering: A framework for representation and analysis. *IEEE Trans. Softw. Eng.*, 34:133–153, Jan. 2008.

[12] M. L. Johnson, S. M. Bellovin, R. W. Reeder, and S. E. Schechter. Laissez-faire file sharing: Access control designed for individuals at the endpoints. In *Proc. NSPW'09*, pages 1–10. ACM, 2009.

[13] R. E. Kasperson, O. Renn, P. Slovic, H. S. Brown, J. Emel, R. Goble, J. X. Kasperson, and S. Ratick. The social amplification of risk: A conceptual framework. *Risk Analysis*, 8(2):177–187, 1988.

[14] K. Kinder-Kurlanda and C. Eder. Under lock and key? Setting up a secure data center at GESIS in Germany. In *39th annual IASSIST conference*, May 2013.

[15] H. R. Lipford and M. E. Zurko. Someone to watch over me. In *Proc. NSPW'12*, pages 67–76. ACM, 2012.

[16] M. S. Lund, B. Solhaug, and K. Stølen. *Model-Driven Risk Analysis - The CORAS Approach*. Springer, 2011.

[17] Managing information security risk: Organization, mission, and information system view. NIST SP 800-39, National Institute of Standards and Technology, Mar. 2011.

[18] E. Ostrom. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge Univ. Press, 1990.

[19] K. Rannenberg. Multilateral security. A concept and examples for balanced security. In *Proc. NSPW'00*, pages 151–162. ACM, 2000.

[20] G. T. Savage, T. W. Nix, C. J. Whitehead, and J. D. Blair. Strategies for assessing and managing organizational stakeholders. *The Executive*, 5(2):61–75, 1991.

[21] N. Schumann and R. Mauer. The GESIS data archive for the social sciences: A widely recognised data archive on its way. *International Journal of Digital Curation*, 8(2):215–222, 2013.

[22] R. T. Simon and M. E. Zurko. Separation of duty in role-based environments. In *Proc. 10th Computer Security Foundations Workshop*, pages 183–194, June 1997.

[23] F. Swiderski and W. Snyder. *Threat Modeling*. Microsoft Press, 2004.

[24] W. Tolone, G.-J. Ahn, T. Pai, and S.-P. Hong. Access control in collaborative systems. *ACM Comput. Surv.*, 37(1):29–41, Mar. 2005.

[25] M. E. Zurko and R. T. Simon. User-centered security. In *Proc. NSPW'96*, pages 27–33. ACM, 1996.