

Asset-Centric Analysis and Visualisation of Attack Trees

Christopher Schmitz¹, André Sekulla², and Sebastian Pape¹

¹ Goethe University Frankfurt, Frankfurt am Main, Germany
{christopher.schmitz; sebastian.pape}@m-chair.de

² University of Siegen, Siegen, Germany
andre.sekulla@uni-siegen.de

Abstract. Attack trees are an established concept in threat and risk analysis. They build the basis for numerous frameworks aiming to determine the risk of attack scenarios or to identify critical attacks or attack paths. However, existing frameworks do not provide systematic analyses on the asset-level like the probability of successful or near-successful attacks on specific assets. But these insights are important to enable decision-makers to make more informed decisions. Therefore, a generic approach is presented that extends classical attack tree approaches by asset-specific analyses. For this purpose, the attack steps in the attack trees are annotated with corresponding assets. This allows identifying the attack paths each asset is exposed to. In combination with the standard attack tree parameter “probability of attack success”, a set of complementary attack success and protection metrics can be applied on each step of the paths. Furthermore, an integrated visualisation scheme is proposed that illustrates the results in a comprehensible way so that decision-makers can intuitively understand what the metrics indicate. It also includes several features improving usability and scalability. As proof of concept, we have implemented a prototype of our proposed method.

Keywords: Attack trees · attack graphs · security metrics · assets · visualisation

1 Introduction

Attack trees are an established concept in threat and risk analysis. Security analysts can use them to determine or compare the risk of attack scenarios, to identify the most likely attack paths or to detect the most serious attack steps. However, current attack tree approaches do not link this information to the asset-level. Thus, they do not provide any systematic information on the assets’ security or risk level. This includes, for example, the probability that an asset will be subject to a (non-) successful attack but also its impact on the overall risk. This information can be of crucial importance from a decision-making perspective, especially when it comes to the question of how to protect against certain attack scenarios in a resource-efficient way which often requires a prioritisation. To address this issue, a generic approach is presented that complements existing attack tree-based risk assessment frameworks, such as the LiSRA framework [20], with meaningful information on the asset-level. This helps to assess the assets’ security level in a systematic and less subjective way and helps decision-makers in mitigating the most critical assets first.

These insights are then visualised in a proper way so that decision-makers can intuitively understand possible attack routes as well as the individual attack chances and the risk sensitivity of each asset. This helps to reflect on the implemented security measures in order to better protect the assets at stake and support the analysis of the underlying risk assessment framework.

The approach consists of the following steps: First, the attack steps in the attack trees are annotated with corresponding assets. This link builds the basis for any security-related analysis with respect to these assets. The attack trees are then transformed into the new concept of “asset-centric attack graphs”. They illustrate the attack paths leading to each individual asset and show which assets have to be attacked in which sequence to perform a successful attack. In combination with the “probability of attack success”, which is a very common parameter for attack trees, new graph metrics are developed that provide meaningful insights on the asset-level [12]. These metrics are developed in such a way that they can be combined in a complementary way. This enables to visualise them in an integrated way so that decision-makers can intuitively understand the assets’ security level. Since real-world infrastructures can be very complex and difficult to comprehend it is also shown how the visualisation scheme can cope with large infrastructures and complex attacks scenarios.

The remainder of this paper is organised as follows. After discussing the background and the related work in Sect. 2, Sect. 3 explains how asset-centric attack graphs can be derived from classical attack trees. Furthermore, meaningful graph metrics are proposed. Sect. 4 then presents the proposed visualisation scheme that describes how all the metrics can be illustrated in an intuitive way. Sect. 5 gives insights on the prototype implementation and evaluates the fulfilment of the requirements that were specified before. Finally, Sect. 6 concludes and points out future work ideas.

2 Background and Related Work

This section gives a brief introduction on attack trees and attack graphs and also gives an overview of related work in both areas.

2.1 Attack Trees

Attack trees are an established concept in threat and risk analysis initially introduced in 1999 by Schneier [21]. An attack tree illustrates an attack scenario from an attacker’s perspective and represents it in a hierarchical, tree-based structure. The attack goal is located in the root node of the tree and is subsequently decomposed into more fine-grained attack steps using logical operations. The leaf nodes finally describe atomic attacker activities. Besides the standard OR and AND operators it is also possible to model more sophisticated operators like “sequential AND” (SAND) described by Jhavar et al. [7]. It is used to consider the sequence of attacker activities.

The principal idea of attack tree-based risk assessment approaches is to analyse the tree nodes with risk-related parameters. Common parameters are the probability of attack success, the costs to perform an attack step or the required skill level. Using bottom-up algorithms these parameters are then propagated up the tree in order to

calculate the values for the entire attack tree [12]. There is a large number of approaches determining the risk of attack scenarios like this, for example, the ADTool or the LiSRA framework [11, 12, 20]. However, only a few approaches consider assets as a central artefact within attack trees [17], and there is no approach that provides systematic analyses of (aggregated) asset classes [12].

2.2 Attack Graphs

An attack graph is an abstraction of attack paths in an infrastructure. The first concept of an attack graph has been proposed by Phillips and Swiler in 1998 [18]. Much work has already been done to define security metrics that can be applied on attack graphs [6, 12, 13, 24]. Several metrics have been defined that are based on the number and the structure of attack paths like the shortest path metric, the number of paths metric, the (normalised) mean of path lengths metric, or the median of path lengths metric [6]. In contrast to these metrics, that only rely on the attack paths, Wang et al. combine path information with the probabilities that specific exploits are executed. In this way, they determine the probability of multi-step attacks [22]. Although this is a specific metric for network security that does not aim to rate or compare certain assets, reachability metrics, in general, build the basis for some of the metrics introduced in the present work. It is shown how reachability metrics can be applied to asset-centric attack graphs. Furthermore, different variations of the standard reachability metric are, amongst other metrics (like the asset risk sensitivity), proposed in this work. The variations indicate, for instance, the probability of near-successfully attack specific assets. The combination of these variations enables an integrated perspective which constitutes one of the central elements of the visualisation scheme. Sawilla and Ou follow a different approach and apply an adapted Google PageRank algorithm on attack graphs in order to rate assets [19]. The assets are associated with vulnerability data from public databases. The algorithm evaluates the assets only in relation to each other. Therefore, it cannot be used to measure the assets' security absolutely. There also exist other approaches that consider assets as central artefacts [3, 8, 9]. However, they cannot be simply applied on (existing) attack trees but they require in-depth knowledge in threat modelling. Apart from that, systematic analyses of (aggregated) asset classes have not been covered by any of these approaches.

Besides that, a key limitation of many attack graph models is that they lack in scalability which makes it challenging to manage their complexity in user interaction. This is mainly due to the fact that real-world attack scenarios can be very complex. As a consequence, many of the attack graph models are too complex to be objectively evaluated by humans in a reasonable time [4, 5, 14, 15, 23]. However, efforts have been made to address this issue. Noel and Jajodia have proposed a technique to collapse attack graph elements through hierarchical aggregation so that attack graphs can be viewed and analysed at different abstraction levels [15]. In 2005, they described a filtering approach that allows the user to filter graph elements so that only the attack subgraphs of interest are shown [14]. Williams et al. represent attack graphs on the basis of treemaps (instead of classical node-link graphs), and they make use of spatial grouping and colour-coding to indicate the level of compromise [23]. Furthermore, they automatically group hosts

with similar levels of compromise. Another approach aiming to reduce the complexity of attack graphs has been presented by Homer et al. [4]. They propose a technique to systematically identify and remove “useless” attack paths. Additionally, their approach performs a grouping of similar attack steps. However, these approaches were not designed to be applied to assets or asset hierarchies. In general, most visualisation approaches in information security are “special-purpose representations” [2].

3 Asset-Centric Analysis of Attack Trees

This section first introduces a scenario that is used to explain the approach. Furthermore, two exemplary attack trees are described for this scenario. Afterwards, it is explained how such trees can be annotated with assets. On this basis, the transformation rules are described to transform attack trees into asset-centric attack graphs. These transformations are necessary for all further asset-specific analyses. Finally, the metrics that build the foundation of the assets’ analyses are presented.

3.1 Scenario Description

The infrastructure depicted in Fig. 1 is adapted from Homer et al. [4]. It illustrates a realistic network infrastructure which consists of three subnets: a demilitarized zone (DMZ), an internal office IT and an industrial control system (ICS) to control critical infrastructure components of an energy provider. Only the DMZ is directly accessible from the Internet. From there, the internal office IT can be accessed through a firewall. Only the Citrix server, which is located in the office IT, has access to the ICS - more precisely to the data historian that again has direct access to the communication server. The goal of the first attack scenario is to gain access to the communication server to modify its control logic. The scenario is illustrated in attack tree T_1 which is depicted in Fig. 2. It is assumed that three attack paths lead to the communication server. For illustration purposes, we give concrete examples for the attacks required to complete the first attack path. For the other paths, we focus more on the actual routes.

An attacker could run a vulnerability scanner from outside the network to discover vulnerabilities on the webserver (a_1). As a result, he might be able to perform a command injection attack on the file server (a_2), and could then access the Citrix server by replacing a binary or shell script on the file server that will be executed on the Citrix server (a_3). From there, the data historian in the ICS could be reached by attacking the remote services (a_8). Due to the direct access to the communication server from the data historian, the attacker could finally execute malicious code on the target (a_9). So the first attack path is:

$$a_1 \rightarrow a_2 \rightarrow a_3 \rightarrow a_8 \rightarrow a_9$$

Alternatively, one could target the VPN server (a_4) and attack the Citrix server from there (a_5). Then, an attacker could perform a_8 and a_9 in the ICS, as described above. Therefore, the second attack path is:

$$a_4 \rightarrow a_5 \rightarrow a_8 \rightarrow a_9$$

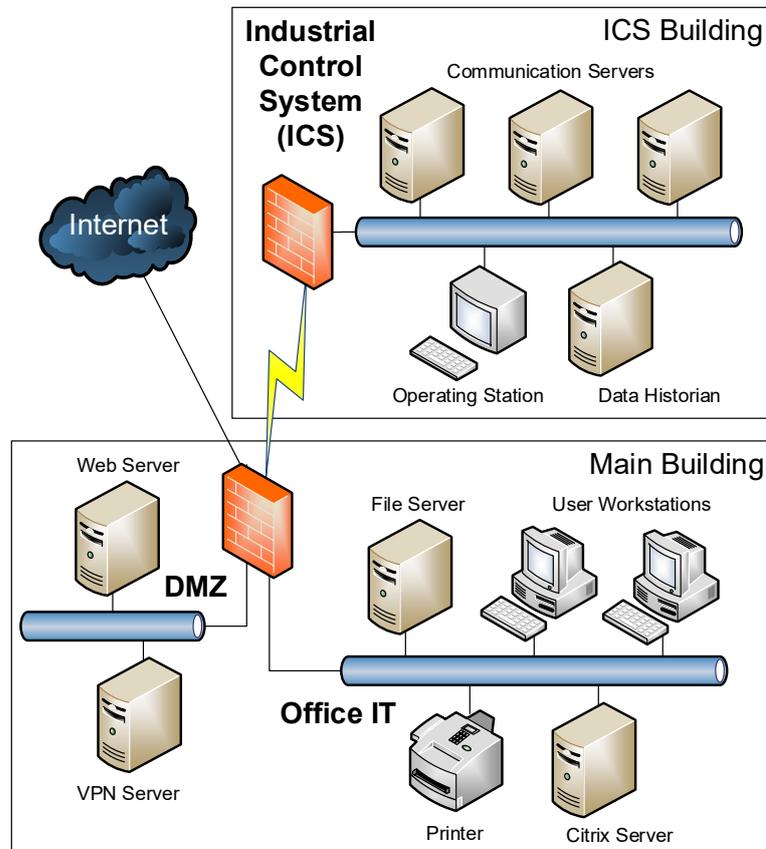


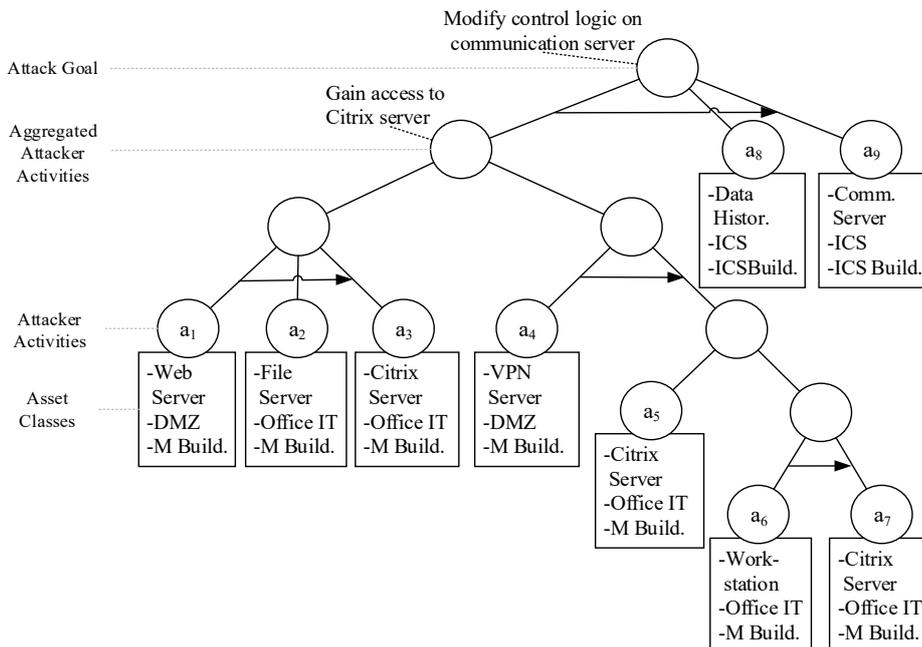
Fig. 1: Exemplary Infrastructure

Instead of directly attacking the Citrix server an alternative attack path is to first attack a workstation:

$$a_4 \rightarrow a_6 \rightarrow a_7 \rightarrow a_8 \rightarrow a_9$$

The workstation may not be as protected as the Citrix server but may have more permissions to access the Citrix server than an external connection via the VPN server allows. It is assumed that the attacker has the privileges for lateral movements in the network unless otherwise specified by defining a specific attack in order to gain access to another network or asset in general.

The second attack scenario is represented by attack tree T_2 which is shown in Fig. 3. It is used to demonstrate the scalability for multiple attack trees. Unlike the first tree, it does not target the communication server but the workstation.

Fig. 2: Attack Tree T_1 with Corresponding Assets

3.2 Annotation of Attack Trees with Assets

The attacker activities of an attack tree can be annotated with different assets, e. g. server assets, security zones or buildings. From a modelling perspective, it makes more sense to annotate asset classes instead of specific assets because of their higher abstraction level. For reasons of readability the terms asset is used in the following instead of “asset class”. A comprehensive overview of assets can be found in the ISO/IEC 27005 risk management standard that comes up with a hierarchical overview of assets. It differentiates between primary assets and supporting assets. Primary assets are the organisation’s business processes & activities and sensitive information, such as trade and business secrets. Both of them can be of crucial importance to an organisation’s success. More interesting from a security engineering perspective are supporting assets because they need to be protected by security measures in the first place. This also becomes clear from the ISO/IEC 27005 definition: “These assets have vulnerabilities that are exploitable by threats aiming to impair the primary assets of the scope (processes and information).” Therefore, only supporting assets are used for annotation.

The standard defines an asset hierarchy covering a wide range of assets, from hardware over software to location assets that include, for instance, security zones and buildings. However, especially for specific attacks, it makes sense to refine these assets with respect to attack-relevant characteristics to cope with attack scenarios targeting more specific types of attack vectors. For example, the asset smart meter (which is relevant for the electric sector) could be differentiated with respect to the supported

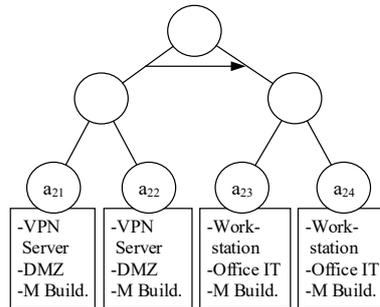


Fig. 3: Attack Tree T_2 with Corresponding Assets

remote data transmission standard (GSM / GPRS, WiFi, Bluetooth, Ethernet etc.). This allows that even specific attacks, such as those targeting only smart meters supporting a WiFi transmission, can be precisely assigned to individual assets.

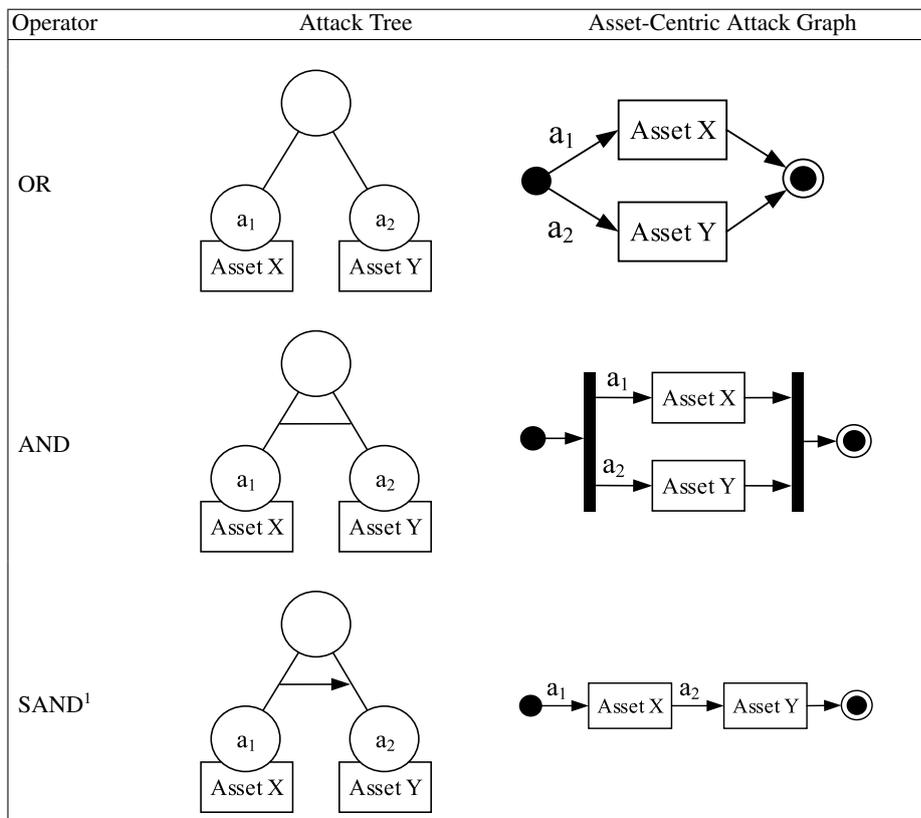
The assigned assets should have the same level of abstraction as the respective attacks. Accordingly, it may be useful to merge similar assets, e. g. comparable workstations that are exposed to similar attacks. A similar host-grouping is applied by Homer et al. in order to reduce complexity [4]. On the other hand, in case of very heterogeneous assets, it makes sense to split these assets into different assets. This enables a more fine-grained analysis. For example, workstations in the ICS might be subject to different attacks than workstations in the office IT.

3.3 Transformation of Attack Trees into Asset-Centric Attack Graphs

Annotated attack trees can be transformed into so-called “asset-centric attack graphs”. The general idea of attack graphs is to systematically illustrate all possible attack paths that are required to achieve the goal of an attack scenario. Asset-centric in this context means that each node represents an asset. Therefore, asset-centric attack graphs show which assets have to be attacked in which sequence to perform a successful attack. Attack paths are thus represented on the level of network topology which typically increases the comprehensibility [4].

The general idea to transform attack trees into attack graphs is already described in literature [10]. The rules to transform an asset-annotated attack tree into an asset-centric attack graph is shown in Fig. 4. If an asset is represented by different nodes in the graph after applying these rules, they must be merged. It is also important to ensure that no cycle is created, as the graphs used in the following are assumed to be acyclic. This can either be ensured by post-annotation checks that indicate which (sub) trees have to be remodelled or by real-time checks which requires a tool-based annotation of assets. After applying the transformation rules on the initial attack trees in Fig. 2 and Fig. 3 the asset-centric attack graph in Fig. 5 is constructed.

The primary purpose of the resulting graph in Fig. 5 is to provide a better understanding of the relation between attacks paths and assets. This is why it concentrates on physical assets (server and workstation assets) only. It must be noted that in case of



¹ Sequential AND

Fig. 4: Transformation of Attack Tree to Attack Graph Patterns (based on [10])

parallel attack steps of different attack paths the tree does not allow to unambiguously distinguish which attack steps belong to which path (e. g. whether a_6 or a_{23} is a direct successor of a_4). This is different for the data structure used for the analysis. However, the graph still shows the sequence of attack steps required to successfully attack a certain asset. This helps to identify neuralgic points, for instance, that all attacks on the communication server run over the Citrix server and the data historian. It also gives a first indication of how exposed the assets are to attacker activities.

The presented approach could generally directly start with adequately constructed attack graphs without considering attack trees before. However, to extend the scope to the popular concept of attack trees, it makes sense to describe the process starting with attack trees.

3.4 Security Metrics

In this subsection, we introduce security metrics that can be applied on the assets in an attack graph. The probability of attack success is a very common metric for attack

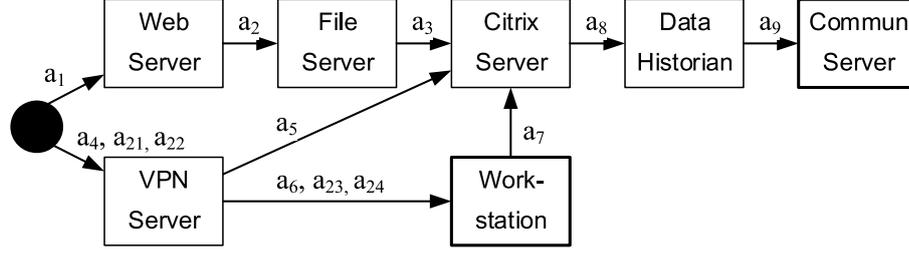


Fig. 5: Asset-Centric Attack Graph Focusing on Physical Assets

trees. The counterpart for the analysis at the asset-level, that is introduced here, is the *probability of successful asset attack*. Another metric is the asset (non-) reachability that describes the probability that an attacker is (not) able to reach the targeted asset at all. Both metrics can be used in combination to derive for each asset how many of the attack attempts can be successfully mitigated. These metrics also allow to derive the probability that an asset is reached but not successfully attacked, in the following referred to as *near-successful asset attack*, although those attacks can even be more challenging than others in practice. Another central metric is the *risk sensitivity* of an asset which measures the maximum impact a specific asset can have on the overall risk.

Successful Asset Attack To determine the attack success for an asset v one must first calculate the probability for each attack path leading up to this asset. Then, one can determine the probability that an attacker can follow at least one of these attack paths until he reaches the targeted asset.

To successfully run through an attack path all of its attacker activities have to be successfully performed. Thus, the success chances for an attack path results from the product of the probabilities of all involved attacker step. Eq. (1) shows the probability to complete path i . $P(\alpha_{ij})$ is the probability to successfully perform attack step j of attack path i . Path i is passed through until the k -th step is reached. k specifies after how many steps an asset v is reached in a particular attack path i . It is returned by the following function: $k_i^v := f(i, v)$. So $P(s_i^v)$ calculates the probability for the complete path i .

$$P(s_i^v) = P\left(\bigcap_{j=1}^{k_i^v} \alpha_{ij}\right) = \prod_{j=1}^{k_i^v} P(\alpha_{ij}) \quad (1)$$

Often an asset can be reached via several attack paths so an attacker can try different ways one after another. The probability to successfully perform (the first k steps of) at least one attack path results from the union set of all attack paths I_v containing asset v . This is reflected by the inclusion-exclusion principle, also known as sieve formula presented in Eq. (2) [1].

$$P(S^v) = P\left(\bigcup_{i=1}^{|I_v|} s_i^v\right) = \sum_{t=1}^{|I_v|} (-1)^{t-1} \left(\sum_{1 \leq i_1 < \dots < i_t \leq |I_v|} P(s_{i_1}^v \cap \dots \cap s_{i_t}^v) \right) \quad (2)$$

It defines a summation over all t -element subsets $\{i_1, \dots, i_t\}$ for each attack path $\{1, \dots, I_v\}$ that contain asset v . It must be noted that the union of attack paths is specified by Eq. (3) ensuring that each attacker activity $\alpha_{ik_i^v}$ is considered only once, even if it appears in several attack paths. The rationale is that a single attacker activity (that is defined from one asset to another) must not be counted multiple times, even if it appears in several paths running in parallel. For example, if the attacker activity a_6 (from the VPN server to the workstation) would occur in both attack trees T_1 and T_2 , it would still only be included once in the calculation of the workstation's attack probability. The reason for this is that it is the same attack step by definition.

$$P(s_{i_1}^v \cap \dots \cap s_{i_t}^v) = \prod_{\alpha_{ik} \in \{s_{i_1}^v, \dots, s_{i_t}^v\}} P(\alpha_{ik_i^v}) \quad (3)$$

Asset Non-Reachability The reachability is calculated similarly to the attack success. The only difference is that an attacker only has to pass the first $k - 1$ attack steps because then he can reach the asset and can continue with $k - th$ step directly targeting the final asset. Therefore, the reachability of an asset v for attack path i is represented by Eq. (4).

$$P(r_i^v) = P\left(\bigcap_{j=1}^{k_i^v-1} \alpha_{ij}\right) = \prod_{j=1}^{k_i^v-1} P(\alpha_{ij}) \quad (4)$$

The probability that at least one of the $|I_v|$ attack paths are performed successfully is calculated using the sieve formula (see Eq. (5)).

$$P(R^v) = \sum_{t=1}^{|I_v|} (-1)^{t-1} \left(\sum_{1 \leq i_1 < \dots < i_t \leq |I_v|} P(r_{i_1}^v \cap \dots \cap r_{i_t}^v) \right) \quad (5)$$

The constraint that each attacker step must not appear multiple times also applies for the reachability as shown in Eq. (6).

$$P(r_{i_1}^v \cap \dots \cap r_{i_t}^v) = \prod_{\alpha_{ik} \in \{r_{i_1}^v, \dots, r_{i_t}^v\}} P(\alpha_{ik_i^v}) \quad (6)$$

The probability of an asset's non-reachability builds the counterpart that indicates how many per cent of all attacks do not reach the asset v . It is calculated as the complementary probability and is shown in Eq. (7).

$$P(\overline{R^v}) = (1 - P(R^v)) \quad (7)$$

Near-Successful Asset Attack The probability that an asset is reached but not successfully attacked is derived from both the attack success and the reachability metric. That means the first $k - 1$ attack steps are successful but the $k - th$ step fails. It is shown in Eq. (8).

$$P(n_i^v) = \prod_{j=1}^{k_i^v-1} P(\alpha_{ij}) \times \prod_{j=k_i^v}^{k_i^v} (1 - P(\alpha_{ij})) \quad (8)$$

$$= \prod_{j=1}^{k_i^v-1} P(\alpha_{ij}) \times (1 - P(\alpha_{ik_i^v})) \quad (9)$$

Finally, $P(n_i^v)$ is inserted into the sieve formula in Eq. (10) with the constraint shown in Eq. (11).

$$P(N^v) = P\left(\bigcup_{i=1}^{|I_v|} n_i^v\right) = \sum_{t=1}^{|I_v|} (-1)^{t-1} \left(\sum_{1 \leq i_1 < \dots < i_t \leq |I_v|} P(n_{i_1}^v \cap \dots \cap n_{i_t}^v) \right) \quad (10)$$

$$P(n_{i_1}^v \cap \dots \cap n_{i_t}^v) = \prod_{\alpha_{ik_i^v} \in \{r_{i_1}^v, \dots, r_{i_t}^v\}} P(\alpha_{ik_i^v}) \quad (11)$$

Additionally, it is possible to derive the asset's self-protection capability from both the *near-successful asset attacks* and the *successful asset attacks*. The self-protection capability is independent of exogenous infrastructural factors, i. e., the self-protection capability of a specific server asset is not influenced by its location in a network topology.

Asset Risk Sensitivity An asset's risk sensitivity determines how sensitive the overall risk reacts on changes to an asset's vulnerability. So it also serves as a measure of protection efficiency. It is assessed as follows: first, all attacker activities directing to the asset to be analysed are identified. They can be read directly from the asset-centric attack graph (see Fig. 5). The Citrix server, for instance, is exposed to the three attacker activities a_3 , a_5 and a_7 . These attacker activities are then assumed to be completely secure respectively completely insecure. The overall risk is then simulated for both states. The difference between both risk values (maximum risk range) represents the maximum impact, *ceteris paribus*, the asset can have on the overall risk. Technically, the insecure state is modelled by temporarily setting the probability of successful asset attack for the k -th attack step required to reach asset v in attack path i to 1. Eq. (12) expresses the state more formally.

$$P_{insecure}^v := P(\alpha_{ik_i^v}) = 1 \quad \forall i \in I_v \quad (12)$$

The secure state is modelled similarly but with a probability of a successful asset attack of 0 (see Eq. (13)).

$$P_{secure}^v := P(\alpha_{ik_i^v}) = 0 \quad \forall i \in I_v \quad (13)$$

The approach presented in this paper extends basic attack tree approaches by asset-specific analyses. Such an attack tree approach can be used to simulate the risk for both states (e. g. LiSRA) [20]. Finally, the difference of both risk values determines the risk sensitivity of an asset v , denoted as ρ^v (see Eq. (14)).

$$\rho^v = \Delta Risk(P_{secure}^v, P_{insecure}^v) \quad (14)$$

In Sect. 4 it is shown how the proposed set of complementary metrics can be visualised in an integrated way.

4 Asset-Centric Visualisation of Attack Graphs

This section proposes a visualisation scheme illustrating the metrics presented in Sect. 3.4 in a comprehensible and self-explanatory way. The scheme aims to support decision-makers in analysing attack scenarios and to enable more informed decisions. As already discussed, real-world attack trees and attack graphs are often far too complex to be objectively evaluated by humans in a reasonable time. For example, Homer et al. have automatically generated the attack graph for basically the same attack scenario also used in this work. They have used the popular MulVAL (Multihost Multistage Vulnerability Analysis) attack graph tool suite [16]. Although the scenario does not seem to be complex, the resulting graph consists of about 130 nodes [4]. This example demonstrates the need for features (cf. Sect. 4.3) supporting decision-makers to analyse complex infrastructures with a large number of assets and attacks. They are proposed in the following.

4.1 Requirements

The authors have elicited the following key requirements and criteria that have to be fulfilled from a practical viewpoint:

- R1** An understandable and comprehensible visualisation of the results. Even users with a lack of technical or security know-how should be able to understand and use the results. This ensures that also less specialised staff from small and medium-sized organisations are able to benefit from the approach. Therefore, the metrics presented must be self-explanatory so that it becomes clear what the metrics indicate.
- R2** The most critical threats and attack targets must be identifiable and stand out from less critical threats and attack targets. The user's focus should be immediately directed to the most critical points of the system in order to be able to take measures as quickly and as effective as possible.
- R3** Besides that, another key requirement is scalability. A scalable visualisation enables to analyse complex real-world scenarios with large attack trees. Although this aspect is of high practical relevance many attack graph approaches have fundamental scalability problems [15].

4.2 Metrics Visualisation

Fig. 6 illustrates an overall view that incorporates the proposed metrics into the asset-centric attack graph.

Attack Success The attack success metrics introduced in Sect. 3.4 have been developed in such a way that they can be combined in a complementary way. For this purpose, the values of the three core metrics are presented in a bullet-based scheme. Each metric is represented by different coloured bullets (red, green, white). Their values are then mapped to the corresponding number of bullets on a scale with 20 bullets, i. e. a probability of successful asset attack of 0.50 yields 10 bullets. The concept of coloured bullets with a scale of 5 per cent per bullet makes it easy to comprehend and to compare the information at first sight. The colour coding for the metrics is shown in Tab. 1.

Table 1: Symbols for Metrics Visualisation

Symbol	Formula	Description
●	$P(S^v)$	Probability that an attacker successfully attacks asset v (successful asset attack)
●	$P(N^v)$	Probability that an attacker reaches but not successfully attacks asset v (near-successful asset attack)
○	$P(\bar{R}^v)$	Probability that an attacker does not reach asset v (asset non-reachability)
◉	ρ^v	Risk sensitivity of asset v

Fig. 6 shows the use in practice. The boxes at the bottom of each asset icon show the integrated view. For example, 50% (10 white bullets) of the attacks do not reach the Citrix server. 40% (8 red bullets) of all attacks are successful, whereas 10% (2 green bullets) can be blocked at the asset. Additionally, the relation between red and green bullets of an asset indicates its self-protection capability. Moreover, it can be seen that the percentage of attacks reaching a certain asset (sum of red and green bullets) decreases the deeper the asset is located in the infrastructure and the less inbound bullets paths it has. The same holds for the percentage of successful attacks per asset (number of red bullets), and therefore also for the number of blocked attacks (number of green bullets).

Risk Sensitivity Another metric is the risk sensitivity of an asset that measures the maximum impact an asset can have on the overall risk. It is illustrated with a coloured protection shield at the right top corner of each asset (from green=*low* to red=*high*). The green shields for the two ICS servers in Fig. 6, for instance, indicate that the overall risk reacts less sensitive to changes of the ICS servers' vulnerability than on the vulnerability of the Citrix server which occurs in all attack paths. Although the adverse impact of a compromised ICS server is much higher, far fewer attacks do reach these servers as indicated by the attack success metrics.

Attack Paths Besides the risk sensitivity and the attack success metrics, it is also essential to understand the attack paths and the attacker activities leading to the respective assets. An attack path connects the assets and illustrates which assets have to be attacked in which sequence. Moreover, the number of inbound and outbound attack paths per asset is illustrated as shown in Fig. 6. They represent to which attacker activities an asset is exposed to and which attacks can be performed from which asset. The number of attacks from asset to asset is represented by the width of the respective edges and is also shown textually. For example, the workstation can be attacked by three different attacker activities - all of them require the VPN server to be successfully attacked first. More detailed information is provided by a mouseover function that displays the concrete attacker activities, their success chances and also refers to the attack trees they originate from.

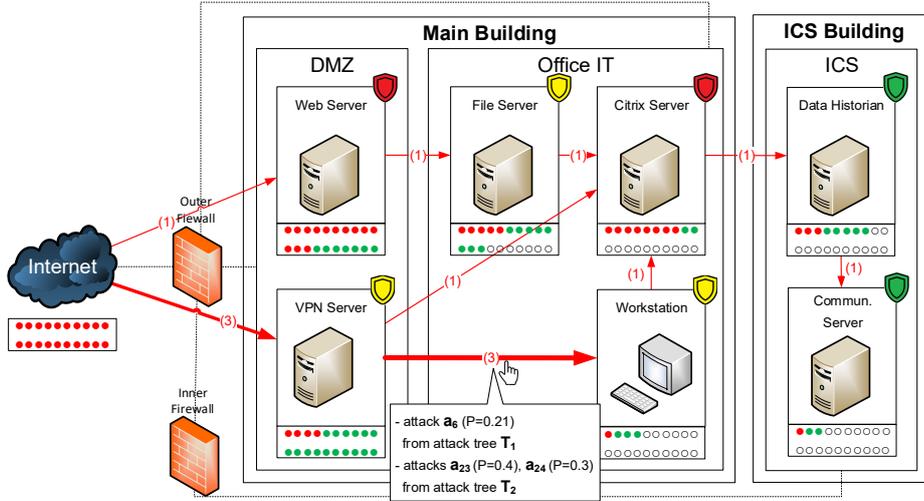


Fig. 6: Detailed View

4.3 Usability and Scalability Features

This section presents features to provide a more usable and scalable view on the security-relevant aspects of the analysed attack scenarios. This is important since it poses a big challenge in practice.

Layered View The overall view depicted in Fig. 6 provides detailed insights for each asset. Although this detailed view can be reasonable not all details and aspects are always necessary. For complex infrastructures, it can even produce an information overload. Therefore, strong scalability features are needed.

Following the ISO/IEC 27005 definition assets can be structured hierarchically so that, for instance, each Exchange server asset is part of the mail server asset which again is part of the general server asset and so on. Asset hierarchies can also be defined individually. For example, each server type is part of a specific security zone, as long as this definition matches the asset annotation in the attack trees. These hierarchies enable to aggregate all metrics to the next higher level (e. g. from server level to level of security zones). This can be achieved by calculating the median and the standard deviation for the lower level assets. Although the median is not an exact metric it still provides a good overview of the asset's security. Fig. 7 shows such a high-level perspective where all metrics at the physical asset level are aggregated to the level of security zones. For the ICS zone this yields to $P(S^{ICS}) = \text{Median}(0.15, 0.05) = 0.10$, $P(U^{ICS}) = \text{Median}(0.25, 0.10) = 0.175$, and $P(R^{ICS}) = \text{Median}(0.6, 0.85) = 0.725$. The presentation of the standard deviations is triggered by a mouseover effect.

Additionally, the number of inbound and outbound attack paths to and from each asset is aggregated to the next higher level by summing. For example, the DMZ has 4 inbound attack paths (1 to the webservice and 3 to the VPN server) and 5 outbound

attack paths leading to the office IT (1 to the file server, 1 to the Citrix server and 3 to the workstation). This is calculated by summing up the paths for the respective asset.

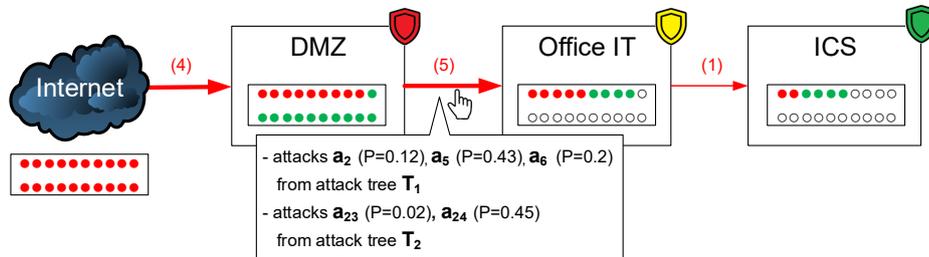


Fig. 7: High-Level View on the Security Zones

Often there is the need to have a closer look at certain aspects of the infrastructure in order to analyse them in detail. This is possible by expanding only these aspects. It is demonstrated in Fig. 8 that shows a high-level perspective for all security zones except for the office IT that can now be analysed in detail.

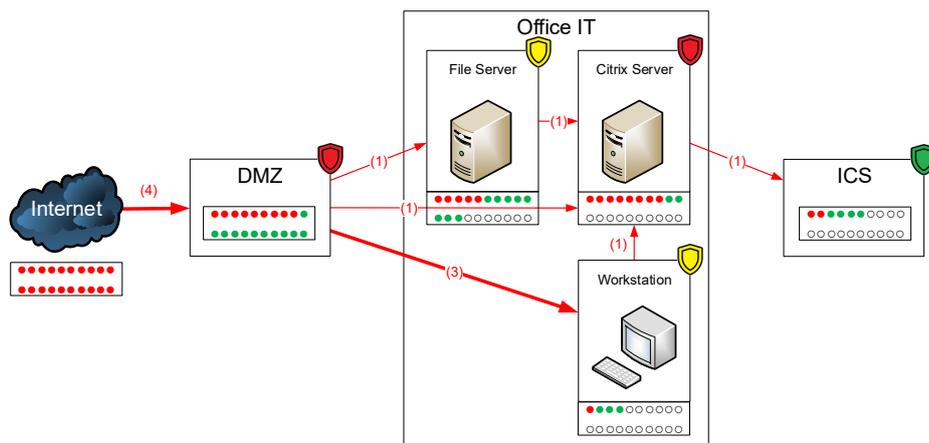


Fig. 8: Focused View on the Office IT

Filter Functionality Also, filtering functions are provided so that the most critical threats stand out from less critical threats and that the user is immediately directed to the neuralgic points. As described before, the visualisation scheme is scalable and supports the analysis of multiple attack scenarios. By default, all scenarios are covered by the analysis. Depending on the target of evaluation it is also possible to display only a subset

of scenarios and assets. For example, if a decision-maker wants to analyse the physical security of the ICS building he can display all physical assets as well as the asset “ICS building” and hide all others.

Another filter enables to only display assets with an actual attack vector. Since no attack path in the presented scenario involves the printer or the operating station both assets are not displayed by default.

Additionally, a threshold can be set for each of the metrics so that only those assets with a score above are shown, i. e. a threshold for the probability of successful asset attack of 0.1 would hide all assets with lower attack chances. This filter can also be used complementary to the layered view. Applied to the view of Fig. 8 the workstation would be hidden in this case. This makes it possible to obtain only the necessary and relevant information which is essential especially for complex infrastructures.

Data Reduction An effective approach for reducing complexity is host-grouping techniques [4]. As described in the section on annotating assets, comparable assets that are exposed to similar attacks should be grouped. For example, the workstation node might represent a grouping of many workstations with comparable configurations. This can significantly reduce the redundancy of data.

5 Prototype Implementation and Evaluation

In this section the developed prototype is described. In addition, it is briefly discussed to what extent the requirements for the visualisation listed above have been met.

5.1 Implementation

The presented approach has also been implemented as a proof of concept in Java. Fig. 9 gives an impression of how the GUI looks like. The tool allows importing a single attack tree or an entire directory of attack trees on the client computer. The import function supports the XML format. The XML structure is based on the widely used ADTool that is also supported by the LiSRA framework [11]. The attacker activities of the imported trees can be annotated with assets directly in the tool. In the lower part of the tool there is an exemplary filtering function to exclude certain assets from the analysis by name or by threshold values. Additionally, important analysis results are summarised on the right side.

5.2 Evaluation of the Visualisation Requirements

The fulfilment of the previously defined requirements is briefly discussed in the following, concentrating on the core aspects. The borders are not clearly demarcated. Some explanations also fulfil another requirement to some extent.

R1 The first requirement addresses aspects like understandability and comprehensibility. According to Homer et al., administrators, who typically deal with network structure plans on a regular basis, will find the representation of attack paths on the level of

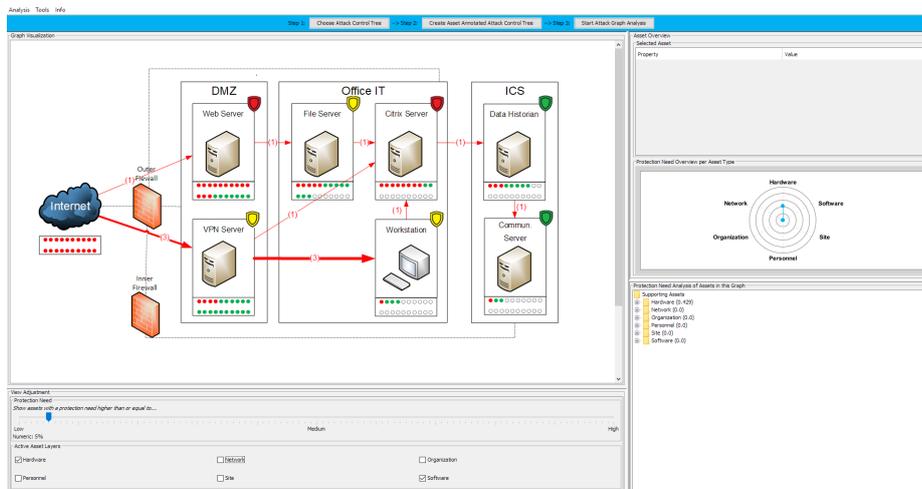


Fig. 9: Prototype Implementation

network topology easier to understand than complex attack graphs [4]. Furthermore, the metrics are presented in an integrated way. Their representation allows a quite intuitive interpretation of the values.

- R2** The most important information must be immediately clear (see R1) and obvious at first glance. This requirement is ensured by the filtering functionality that hides all non-critical assets and attacks. Furthermore, the proposed metrics are displayed graphically next to the corresponding assets. This allows the metrics to be immediately associated with the correct assets.
- R3** To meet the requirement of scalability various features have been integrated. The most important is the layered view. In addition to a very fine-granular view also aggregated views are supported enabling a high-level analysis, as shown in the example of the view for analysing security zones. All metrics are automatically re-calculated according to the chosen abstraction level. Starting from this view, it is possible to navigate through the infrastructure, for instance, by following the critical assets. Additionally, this can be complemented by the filter functionality that allows hiding the non-relevant assets that are out-of-scope or non-critical.

6 Conclusion and Future Work

Attack trees are an established concept in threat and risk analysis that is used to analyse entire attack scenarios and their attacker activities. They enable to identify the most likely attack paths or the most serious attacks. However, there are no attack tree approaches that provide systematic analyses on the asset-level. But this is important from a defenders' perspective to better understand the security and risk of each asset.

Therefore, a novel approach has been proposed that extends attack tree frameworks by linking the attacker activities in the trees to the asset-level. These annotated attack

trees can then be transformed into an asset-centric attack graph that illustrates the attack paths for each asset. Together with the standard attack tree parameter “probability of attack success” these paths enable to apply a set of complementary security metrics that give meaningful insights for each asset. To enable these calculations novel metrics like the asset risk sensitivity or the near-successful asset attack have been proposed, partially based on existing reachability metrics. One of the main contributions concerning the security metrics lies in the combination of (partially new) metrics which enable an integrated visualisation. Furthermore, a visualisation scheme has been developed to integrate these complementary metrics into the asset-centric attack graph. All results are presented in such a way that decision-makers can intuitively understand the attack paths as well as the rationale for the individual attack chances and the risk sensitivity of each asset. Since usability and scalability issues pose a big challenge in the visualisation of attack graphs several features have been proposed to cope with complex attack scenarios. The approach has also been implemented in a prototype. The next step will be to conduct a user study to systematically evaluate and improve the approach from a users’ perspective.

Acknowledgments

This work was partially supported by European Union’s Horizon 2020 research and innovation program from the project CyberSec4Europe (grant agreement number: 830929). We also thank Niklas Paul for his contribution to the prototype implementation.

References

1. Inclusion-exclusion principle. <https://mathworld.wolfram.com/Inclusion-ExclusionPrinciple.html> (2020), accessed: 04 May 2020
2. Fink, G.A., North, C.L., Endert, A., Rose, S.: Visualizing cyber security: Usable workspaces. In: 2009 6th International Workshop on Visualization for Cyber Security. pp. 45–56 (2009)
3. Holm, H., Shahzad, K., Buschle, M., Ekstedt, M.: P² cysemol: Predictive, probabilistic cyber security modeling language. *IEEE Transactions on Dependable and Secure Computing* 12(6), 626–639 (2015)
4. Homer, J., Varikuti, A., Ou, X., McQueen, M.A.: Improving attack graph visualization through data reduction and attack grouping. In: International Workshop on Visualization for Computer Security. pp. 68–79. Springer (2008)
5. Hong, J.B., Kim, D.S., Chung, C.J., Huang, D.: A survey on the usability and practical applications of graphical security models. *Computer Science Review* 26, 1–16 (2017)
6. Idika, N., Bhargava, B.: Extending attack graph-based security metrics and aggregating their application. *IEEE Transactions on dependable and secure computing* 9(1), 75–85 (2010)
7. Jhavar, R., Kordy, B., Mauw, S., Radomirović, S., Trujillo-Rasua, R.: Attack trees with sequential conjunction. In: IFIP International Information Security and Privacy Conference. pp. 339–353. Springer (2015)
8. Johnson, P., Vernotte, A., Ekstedt, M., Lagerström, R.: pwnpr3d: An attack-graph-driven probabilistic threat-modeling approach. In: 2016 11th International Conference on Availability, Reliability and Security (ARES). pp. 278–283 (2016)

9. Johnson, P., Lagerström, R., Ekstedt, M.: A meta language for threat modeling and attack simulations. In: Proceedings of the 13th International Conference on Availability, Reliability and Security. ARES 2018, Association for Computing Machinery, New York, NY, USA (2018), <https://doi.org/10.1145/3230833.3232799>
10. Karray, K., Danger, J.L., Guilley, S., Elaabid, M.A.: Attack tree construction and its application to the connected vehicle. In: Cyber-Physical Systems Security, pp. 175–190. Springer (2018)
11. Kordy, B., Kordy, P., Mauw, S., Schweitzer, P.: Adtool: Security analysis with attack–defense trees. In: Joshi, K., Siegle, M., Stoelinga, M., D’Argenio, P.R. (eds.) Quantitative Evaluation of Systems. pp. 173–176. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)
12. Kordy, B., Piètre-Cambacédès, L., Schweitzer, P.: Dag-based attack and defense modeling: Don’t miss the forest for the attack trees. Tech. rep. (2014)
13. Lippmann, R.P., Ingols, K.W.: An annotated review of past papers on attack graphs. Tech. rep., Massachusetts Inst of Tech Lexington Lincoln Lab (2005)
14. Noel, S., Jacobs, M., Kalapa, P., Jajodia, S.: Multiple coordinated views for network attack graphs. In: IEEE Workshop on Visualization for Computer Security, 2005.(VizSEC 05). pp. 99–106. IEEE (2005)
15. Noel, S., Jajodia, S.: Managing attack graph complexity through visual hierarchical aggregation. In: Proceedings of the 2004 ACM workshop on Visualization and data mining for computer security. pp. 109–118 (2004)
16. Ou, X., Govindavajhala, S., Appel, A.W.: Mulval: A logic-based network security analyzer. In: USENIX security symposium. vol. 8, pp. 113–128. Baltimore, MD (2005)
17. Paul, S., Vignon-Davillier, R.: Unifying traditional risk assessment approaches with attack trees. *Journal of Information Security and Applications* 19(3), 165 – 181 (2014), <http://www.sciencedirect.com/science/article/pii/S2214212614000180>
18. Phillips, C., Swiler, L.P.: A graph-based system for network-vulnerability analysis. In: Proceedings of the 1998 workshop on New security paradigms. pp. 71–79 (1998)
19. Sawilla, R.E., Ou, X.: Identifying critical attack assets in dependency attack graphs. In: European Symposium on Research in Computer Security. pp. 18–34. Springer (2008)
20. Schmitz, C., Pape, S.: Lisra: lightweight security risk assessment for decision support in information security. *Computers & Security* 90, 101656 (2020)
21. Schneier, B.: Attack trees. *Dr. Dobb’s journal* 24(12), 21–29 (1999)
22. Wang, L., Islam, T., Long, T., Singhal, A., Jajodia, S.: An attack graph-based probabilistic security metric. In: Atluri, V. (ed.) *Data and Applications Security XXII*. pp. 283–296. Springer Berlin Heidelberg, Berlin, Heidelberg (2008)
23. Williams, L., Lippmann, R., Ingols, K.: An interactive attack graph cascade and reachability display. In: *VizSEC 2007*, pp. 221–236. Springer (2008)
24. Yusuf, S.E., Hong, J.B., Ge, M., Kim, D.S.: Composite metrics for network security analysis. *Software Networking* 2017(1), 137–160 (2018)