

Intrusion response cost assessment methodology

Chris Strasburg
Department of CS
Iowa State University
cstras@cs.iastate.edu

Natalia Stakhanova
Faculty of CS
University of New Brunswick
natalia@unb.ca

Samik Basu
Department of CS
Iowa State University
sbasu@cs.iastate.edu

Johnny S. Wong
Department of CS
Iowa State University
wong@cs.iastate.edu

ABSTRACT

In this paper we present a structured methodology for evaluating cost of responses based on three factors: the *response operational cost* associated with the daily maintenance of the response, the *response goodness* that measures the applicability of the selected response for a detected intrusion and the *response impact on the system* that refers to the possible response effect on the system functionality. The proposed approach provides consistent basis for response evaluation across different systems while incorporating security policy and properties of specific system environment.

Categories and Subject Descriptors

C.2.0 [Computer-Communication Networks]: Security and protection

General Terms

Security

Keywords

intrusion response assessment, cost-sensitive intrusion response

1. INTRODUCTION

In recent years the trend toward cost-sensitive modeling of response selection became more apparent [3, 1, 2, 6, 4]. The primary aim for applying such models is to balance intrusion damage and response cost to ensure adequate response without sacrificing the normal functionality of the system under attack. However, one of the challenges in applying this approach is defining accurate and consistent measurement of the cost factors on the basis of requirements and policies of the system being protected against intrusions.

One of the primary problems in this context is to identify whether or not a response should be deployed, in other words, what is the best suited action when an intrusion is detected. This problem primarily stems from the fact that

though responses are deployed with the goal of countering an intrusion, they may not only fail but can also lead to undesired effects on the system. Thus, often the primary criteria in response selection mechanisms are the expected effectiveness of the response against the intrusion and its potential negative impact on the system.

The effectiveness of a response refers to both the ability of the response to prevent or mitigate damage from the intrusion and the coverage of the response, i.e., the number of intrusions it can potentially address. One of the intuitive ways to measure the effect of the response is to consider the system resources affected by the intrusion and protected by the response.

Another factor characterizing the response is its potential effect on the system. While the responses are deployed against a detected intrusion, they often alter the state of the system negatively affecting system resources and leading to damage.

Although, the response effect on the detected intrusion and its impact on the system resources are the primary characteristics considered for intrusion response, several important factors remain behind the scene: administrator time and additional system resources (i.e., storage, network bandwidth, etc.) required for response setup and processing. Collectively, we will refer to these factors as operational cost. While this cost does not directly affect the attacked system or the intrusion, it can significantly contribute to the decision of which response to deploy.

In light of the above, we present a structured system independent methodology for the evaluation of responses' cost based on the three parameters: (a) *the response goodness in addressing the detected intrusion(s)* which includes the effectiveness of the response and its coverage capability, (b) *the damage incurred by a response on the system* and (c) *the operational cost of a response on a given system*.

Within this methodology, we propose to assess response impact with respect to resources of the affected system. Our model takes into account the relative importance of the system resources determined through the review of the system policy goals according to the following categories: *confidentiality*, *availability* and *integrity*.

This methodology does not substitute the response selection process in case of detected intrusion, but rather allows

- | |
|---|
| <p>1: The system classification:
 <i>-identify the type of the system according to the security goals</i></p> <p>2: The system policy goals:
 <i>-assign weights to system policy goals for the system</i></p> <p>3: The system resources:
 <i>-enumerate resources available on the given system</i>
 <i>-determine the resource importance for each system policy goal</i>
 <i>-compute the overall resource weight for the system policy</i></p> <p>4: The response taxonomy:
 <i>-identify the responses suitable for the system</i></p> <p>5: The response operational cost:
 <i>-assess the operational cost of the responses</i></p> <p>6: The response goodness:
 <i>-assess the goodness of the responses</i></p> <p>7: The response impact on the system:
 <i>-compute the impact of the available responses</i></p> |
|---|

Figure 1: The methodology for intrusion response cost evaluation.

to evaluate the available responses in a consistent fashion. We have implemented the proposed response cost evaluation methodology and believe that it can be employed to guide system administrators during the response selection process¹.

The main contributions of this work can be summarized as follows:

1. **Structured and comprehensive methodology for assigning response costs:** the proposed model presents the effective roadmap for defining a standardized metric for response cost evaluation.
2. **System independent evaluation model:** the proposed model is adaptable to different environment settings, ie. systems with widely varying operational requirements.
3. **Consistent response metrics:** the proposed evaluation metrics are defined in terms of the system resources that bring a common ground to the assessment process.
4. **Adaptable evaluation metrics:** the response metrics are quantified with respect to the security policies and properties of the specific system. Thus, the computed costs can be effortlessly adjusted as and when the system requirements are modified.

2. RESPONSE COST EVALUATION MODEL

The evaluation of the intrusion response cost is performed in three dimensions: *the operational cost (OC)* of a response in a given environment, that measures various aspects of the response associated with its daily maintenance; *the response goodness (RG)* with respect to detected intrusion(s) that provides a measure of the ability of the corresponding response to mitigate damage caused by the intrusion to the system resources; and finally, *the response impact on the system (RSI)* that quantifies the negative effect of the response on the system resources and that is estimated independently from the response success or failure in countering the intrusion(s).

Intuitively, the combination of the OC and the RSI constitutes the penalty associated with the response, while the RG is the benefit of this response measure. One simple cost model describing the overall measure of response cost RC is:

$$RC = OC + RSI - RG \quad (1)$$

¹The experimental results can be found in [5]

Figure 1 presents the overview of steps for evaluating the response cost RC (following Equation 1). In the proceeding discussion we introduce each step in detail.

Step 1: System classification. The first step in quantifying the cost of a response involves determining the characteristics of the computing environment where the response will be deployed. For example, the system can be classified in terms of the security goals as an *open-access system* with minor or none security restrictions (e.g., public networks provided at airport) or as a *safety-critical system* with emphasis on the service availability. The classification process should provide important insights to the risks that each class of systems can tolerate, and therefore help in measuring the cost for various types of intrusion damages.

Step 2: The system policy goals. The determination of the importance of the system policy goals, and subsequently, the assessment of the potential risks are the responsibilities of the organization to which the system belongs. It is usually a manual process consisting of an informal series of questions such as “Will data be exposed?”, “How critical is the confidentiality of the data?”, “How concerned are we with data integrity?”, “Will service availability be impacted?”, etc. This provides an ad-hoc relative assessment of the system goals for the organization. Based on the above observation, system policy goals can be defined in terms of: *a) Confidentiality* that refers to the imposed restrictions on information flows, e.g., restricted access to data. *b) Integrity* that is a guarantee of the consistency and accuracy of the information or the system computing environment as a whole. *c) Availability* that indicates the requirement of (functionality, storage etc.) service and information availability upon request.

These categories of system goals are ranked according to their importance (a value between 0 for *no importance* and 1 for *absolute importance*) in a particular system type (safety-critical, security-critical, etc.). These decisions can be based on monetary values or other established business metrics for the cost of failure to meet system goals (e.g., the estimated dollar cost of a confidentiality breach). In the case of a classified data processing system (a security-critical system), for instance, data confidentiality may be a 1, indicating the absolute importance of this security facet for this system.

Step 3: System resources. Responses are reactions to the intrusions and are directed to protect the system resources threatened by an attack. System resources can be broadly viewed as the system assets (e.g., host, network, etc.), services provided by the system (e.g., FTP, HTTP, file system, etc.) and users served by the system.

One of the initial steps in the process for computing a response impact measure is the enumeration of the resources available in the considered system. The importance of a resource depends on the system policy goals which in turn depend on the type of system. For example, for a simple Web server, availability is an important policy goal and accordingly important resources will include HTTP. Therefore,

the resources are assigned weights according to their importance for each system policy goal for a specific system. The overall weight of the system resource, denoted by W_{SR} , is computed as a combination of the resource importance for each policy goal category $SR_{importance_i}$ (i is the policy category index) and the system specific category weight $PolicyCategoryWeight_i$ (weight of the i -th policy category index):

$$W_{SR} = \sum_i [SR_{importance_i} \times PolicyCategoryWeight_i] \quad (2)$$

To illustrate this process, let's consider the example of the network interface resource and its importance for each policy category for a public Web server:

Policy Category	Weight	$SR_{importance_{NetworkAccess}}$
Data confidentiality	0	0.1
Data availability	1.0	1.0
Data integrity	0.7	0.1
$W_{NetworkAccess} = 0 \times 0.1 + 1.0 \times 1.0 + 0.7 \times 0.1 = 1.07$		

Step 4: Taxonomy of responses. Once the system goals are identified, its resources are enumerated and their importance is quantified based on the system goals, the next step is to identify the set of responses that are suitable for a system. Generally, the responses are deployed to either counter possible attacks and defend the system resources or regain secure system state. Thus, the selection of applicable responses primarily depends on the identified system resources.

Step 5: Assessment of response operational cost. The assessment of *response operational cost* is generally independent from the system policy and includes the cost for the setup and deployment of the response, and data processing overhead needed to analyze the result of response. For example, “the system logging” response is fairly easy to setup. However, it requires significant storage resources and often incurs high processing overhead. Broadly, the involved operational expenses can be classified on the basis of three requirements: *human resources* which refer to administrator time, *system resources*, which include storage, network bandwidth, processor time, etc., and *direct expenses* which include data processing fees by a third party, subscription service fees, cost of components replacement, etc. Determining these factors is a manual process that involves expert knowledge and a high degree of judgment.

Step 6: Assessment of response goodness. The evaluation of the response goodness includes *a) Selection of the applicable responses for intrusions* and *b) Computation of the response goodness measure in terms of known attacks*

Often the detection mechanism of the intrusion detection system (IDS) provides administrator with a set of alerts indicating potential attacks rather than a specific intrusion. When this situation arises, the response needs to be deployed preemptively on the basis of high likelihood of possible intrusions. In these cases, the response is evaluated based on the number of possible intrusions it can potentially address, and consequently, the number of resources that can be protected

by the response. In practice, the applicability of responses to potential attacks can be determined through the analysis of the existing intrusion signatures in the IDS.

The assessment of response goodness includes a review of the availability of the system resources involved in the intrusion. For example, an alert triggered on TFTP traffic on port 69 is accounted for in the response goodness assessment only if TFTP protocol is currently supported.

The goodness of the response R_i where $i \in [1 \dots m]$ (m different responses) against the intrusion I_j potentially affecting n system resources $SR_1^j, SR_2^j, \dots, SR_n^j$ is computed as follows:

$$RG_{R_i}(I_j) = \sum_{k \in [1 \dots n]} Avail(SR_k^j) \times W_{SR_k} \quad (3)$$

where $Avail(SR_k^j)$ is a binary value that denotes the availability of k -th system resource that can be affected by I_j and W_{SR_k} is the resource weight (as computed by Equation 2). To ensure the consistency of the computed metric, RG values are normalized within a range of $[0, 1]$ by dividing individual $RG_{R_i}(I_j)$ by the normalization term, $MAX(RG(I_j))$ which is the maximum RG value computed for available responses for the intrusion I_j , i.e.,

$$MAX(RG(I_j)) = RG_{R_i}(I_j) \text{ such that}$$

$$l \in [1 \dots m] \wedge \forall i \in [1 \dots m] : RG_{R_i}(I_j) \leq RG_{R_l}(I_j)$$

In the rest of the paper, we will refer to $RG_{R_i}(I_j)$ to mean its normalized valuation.

Step 6: Assessment of the response impact on the system. The impact of a response is evaluated based on the defined system goals and their importance. The impact assessment process for a specific response includes three steps. *First*, identify the system resources affected by each response. *Second*, for each resource, order the responses on the basis of how they are affecting the resource. *Finally*, compute the negative impact of the responses on the associated resource using the ordering obtained above. Eventually, the impact of a response on the system as a whole will be an aggregation of the response's impact on the resources present in the system.

For each response we determine the system resources it may affect. For instance, *blocking a specific subnet* can protect the network interface resource and also disrupt legitimate user activities. After all responses are categorized within the considered system resource, we independently evaluate each system resource. Specifically, all responses affecting the resource are ordered or ranked based on their relative impact on the considered resource, from the greatest impact to the least impact, and assigned an index $i \in [0 \dots (m - 1)]$, where m is the total number of responses in the list corresponding to a particular resource. A response with rank i has more impact on the corresponding resource than the response with rank j ($i < j$). These ranks are based on historical data and/or the expertise of the system administrator. We quantify the impact using the rank as follows:

$$Impact_{R_i, SR} = 1 - \frac{i}{m} \quad (4)$$

where R_i is the i -th ranked response. The resultant valuation is between $\frac{1}{m}$ and 1. To illustrate this process, let's consider the example of the *network interface* resource. The available responses are ranked according to their impact and the corresponding impact quantification is computed as follows:

Rank i	Responses for SR (R_i)	Impact $_{R_i,SR}$
0.	Complete network isolation	$1 - 0/5 = 1.0$
1.	Network isolation: block subnet	0.8
2.	Terminate process	0.6
3.	Delay suspicious process	0.4
4.	Deploy intrusion analysis tools	0.2

Generally, the values determined as a result of ranking are dependent on the characteristics of system environment. As such, changes in the environment, i.e., modifications in the software usage, addition of network equipment, new knowledge or skills gained by the administrator, etc., can affect the order and relative severity of the responses. Thus, as the settings of the environment change, these values may be manually adjusted to more accurately reflect relative damage on the system resources.

The overall impact of the response measure is estimated based on the weight of the system resource for a specific system policy (Equation 2) and the impact value of the response for that resource (Equation 4). The overall rating of the response R_i on the system, the response system impact, denoted by RSI_{R_i} , is computed as follows:

$$RSI_{R_i} = \sum_{SR} \text{Impact}_{R_i,SR} \times W_{SR} \quad (5)$$

Similar to RG valuations (Equation 3), we normalize RSI_{R_i} using the maximum valuation of RSI for any response.

While manual assignment of some values is inevitable, these abstractions allow an expert to focus separately on the technical nature of the responses and the high-level goals of the system. In many cases, two different individuals or groups are uniquely qualified to make the respective technical and policy based decisions. As such if the environment changes, the system administrator can modify the high-level system goals while a technical specialist adjusts response damage factors based on changes to the system or network environment. Such separation of concern reduces the decision complexity, and therefore, the risk of human error.

3. PRACTICAL EXERCISE

To evaluate the practical value of our approach, we conducted an experiment where we asked system administrators to rank the set of response actions using their traditional methods according to responses' priority to be deployed on the system in the case of an SQL injection attack. In the experiment we offered four types of system: *public web server*, *classified research system*, *medical data repository* and *receptionist workstation*. We recruited 9 system administrators with different level of expertise (5 experts and 4 with moderate level of expertise). The motivation for the experiment was to evaluate the consistency of the response cost assessment using our methodology in comparison with the traditional approach primarily based on the manual selection of responses according to the administrator expertise.

Surprisingly, the results showed a substantial variability in the response ranking among administrators. The rank order correlation coefficients between any two rankings are in the range of $\{-0.74, 0.15\}$. This means that ranking is not consistent neither among experts, nor among administrators with moderate expertise level, and consequently, varies from the ranking determined by our approach. As one of the responders noted, the response ranking provided by our method characterized a smooth process for system administrators to follow during an attack, while his personal response preference is an overreaction to the situation.

This provides strong testimony that even experienced administrators need a standardized metric for evaluating intrusion responses that would allow to assess the costs involved in each response deployment in a consistent manner. Our approach can be employed by system administrators to guide them through the response selection process

4. CONCLUSION AND FUTURE WORK

In this paper we have presented a comprehensive and structured methodology for evaluation of response cost. The proposed model identifies three main components that constitute response cost, namely, response operational cost, the response goodness in mitigating the damage incurred by the detected intrusion(s) and the response impact on the system. These response metrics provide a consistent basis for evaluation across systems, while allowing the response cost to be adapted with respect to the security policy and properties of specific system environment. This approach takes advantage of the accuracy inherent in expert assignment of values, and combines it with a structured calculation of relative values, resulting in flexibility and consistency. Importantly, this approach is practically implementable in a real-world environment, making response cost assessment accessible to system administrators with a range of system expertise.

5. REFERENCES

- [1] I. Balepin, S. Maltsev, J. Rowe, and K. Levitt. Using specification-based intrusion detection for automated response. In *Proceedings of RAID*, 2003.
- [2] B. Foo, Y.-S. Wu, Y.-C. Mao, S. Bagchi, and E. H. Spafford. ADEPTS: Adaptive intrusion response using attack graphs in an e-commerce environment. In *Proceedings of DSN*, pages 508–517, 2005.
- [3] W. Lee, W. Fan, M. Miller, S. J. Stolfo, and E. Zadok. Toward cost-sensitive modeling for intrusion detection and response. *J. Comput. Secur.*, 10(1-2):5–22, 2002.
- [4] N. Stakhanova, S. Basu, and J. Wong. A cost-sensitive model for preemptive intrusion response systems. In *Proceedings of AINA*, pages 428–435, Washington, DC, USA, 2007. IEEE Computer Society.
- [5] C. Strasburg, N. Stakhanova, S. Basu, and J. Wong. The methodology for evaluating response cost for intrusion response systems. Technical Report 08-12, Iowa State University, 2008.
- [6] Y.-S. Wu, B. Foo, Y.-C. Mao, S. Bagchi, and E. Spafford. Automated adaptive intrusion containment in systems of interacting services. In *To appear in Journal of Computer Networks*, 2007.