The Challenges of Using an Intrusion Detection System: Is It Worth the Effort?

Rodrigo Werlinger, Kirstie Hawkey, Kasia Muldner, Pooya Jaferian, Konstantin Beznosov

University of British Columbia, Vancouver, Canada
{rodrigow, hawkey, kmuldner, pooya, beznosov}@ece.ubc.ca

ABSTRACT

An intrusion detection system (IDS) can be a key component of security incident response within organizations. Traditionally, intrusion detection research has focused on improving the accuracy of IDSs, but recent work has recognized the need to support the security practitioners who receive the IDS alarms and investigate suspected incidents. To examine the challenges associated with deploying and maintaining an IDS, we analyzed 9 interviews with IT security practitioners who have worked with IDSs and performed participatory observations in an organization deploying a network IDS. We had three main research questions: (1) What do security practitioners expect from an IDS?; (2) What difficulties do they encounter when installing and configuring an IDS?; and (3) How can the usability of an IDS be improved? Our analysis reveals both positive and negative perceptions that security practitioners have for IDSs, as well as several issues encountered during the initial stages of IDS deployment. In particular, practitioners found it difficult to decide where to place the IDS and how to best configure it for use within a distributed environment with multiple stakeholders. We provide recommendations for tool support to help mitigate these challenges and reduce the effort of introducing an IDS within an organization.

Categories and Subject Descriptors

K.6.5 [**Management of Computing and Information Systems**]: Security and Protection; H.5.3 [**Information Interfaces and Presentation**]: Group and Organization Interfaces—*Collaborative Computing*

General Terms

Human Factors, Security Management, Design

Keywords

Intrusion Detection, Usable Security, Collaboration, Qualitative Research, Security Tools, Organizational Factors

Symposium On Usable Privacy and Security (SOUPS) 2008, July 23-25, 2008, Pittsburgh, PA, USA.

1. INTRODUCTION

Security incident response is one key aspect of maintaining organizational security [21]. A critical task during security incident response is detecting that an incident has occurred. Detection may occur through reports from end-users and other stakeholders in the organization, through detection analysis performed on an ad-hoc basis (e.g., hand-crafted scripts that detect anomalies in server logs), or it may be accomplished by using an intrusion detection system (IDS). In general, an IDS monitors and records events in a computer system, performs analysis to determine if the events are security incidents, alerts security practitioners of potential threats, and produces event reports [31]. If the IDS also includes mechanisms to block detected intrusions from entering the organizational infrastructure, it is referred to as an intrusion prevention system (IPS). Security practitioners interact with the IDS through a console, which may be used to either perform administrative functions, such as configuration of sensors, and/or to support event monitoring and analysis. Some of the most popular IDSs include Snort [33], OSSEC HIDS [27], BASE [4], Sguil [32], and Bro [6].

Intrusion detection (ID) is a challenging endeavor, requiring security practitioners to have a high level of security expertise and knowledge of their systems and organization [31, 12]. Traditionally, ID research has focused on technological solutions for improving the accuracy of IDSs (e.g., [8, 18]). Although this is still an active area of research, recent work has also recognized the need to address the human side of ID work (e.g., [12, 22, 36]). This recognition is driven by the fact that while IDSs automate some aspects of the process, human intervention is very much still required. For instance, although an IDS automatically recognizes potential security threats and generates alerts, the alerts need to be analyzed by a human expert, since many are false positives (as many as 99 percent [19]).

From a usability perspective, much of the research has focused on providing visualizations during the monitoring and analysis phases (e.g., [24]), with some claiming these phases to be the most cognitively challenging [35]. However, the initial deployment and configuration of the IDS can also be a barrier to its use. The first author has experienced this first-hand while working as a security consultant at a large telecommunications company from 2002 to 2006. This organization's security team wanted to employ an IDS to improve the organization's security, but had two main concerns about incorporating such a system: (1) Were they going to be able to maintain it? (to ease this burden they had the option of outsourcing the network monitoring, but did not

want to disclose the log files), and (2) Were they going to learn valuable information from the reports (e.g., were there attacks on their systems that needed to be addressed)? Despite assistance from an external company with the initial configuration of the IDS, the security team was unable to customize it and tune it appropriately for the network they were monitoring within a reasonable time frame.

In this paper we report on the challenges of using an IDS, with a particular focus on the initial stages of deployment (i.e., decision making, installation, and configuration). Our motivation for this research arose from the first author's prior industry experience as described above. We also noted that other practitioners had similar difficulties with IDSs through our research conducted for the HOT Admin project, which is investigating the human, organizational, and technological factors that influence security management within organizations (see [15] for an overview and [5], [11], [16], [37], [38] for results to date).

Our findings are based on analysis of nine of the HOT Admin interviews that we conducted with security practitioners, as well as participatory observation in a large academic organization that is in the process of installing an IDS. This rich set of data has allowed us to identify and describe some of the challenges that impact the ability of security practitioners to successfully deploy and maintain an IDS within an organization. These challenges include deciding on the purpose of the IDS, integrating the IDS in the network, working within a distributed environment, and balancing the tradeoff between limiting the number of false positives to achieve usability of the system, while keeping false negatives at a minimum. While some of these challenges may not have obvious solutions, it is important that security practitioners, researchers, and tool developers are aware of the complexity of the full process of deploying an IDS.

Our work has two key contributions. First, we add to the community's understanding of the factors influencing IDS usability. In particular, while prior work has focused on the challenges associated with the monitoring and analysis phases of IDS work, suggesting that these phases are the most cognitively demanding, our results show that the deployment phase also involves challenges, and that these may be significant enough to hinder the very adoption of an IDS within an organization. Second, we provide recommendations and guidelines for mitigating some of the challenges we identify through better tool support.

The remainder of the paper is organized as follows. We begin by presenting the related work in section 2 and our methodology in section 3. In section 4, we describe the IDS tool used during participatory observation, and then present our results related to IDS usability in section 5. We discuss our findings in section 6 before presenting conclusions and future work.

2. RELATED WORK

Before devising support for the human analysts who work with IDSs, it is important to have an understanding of what is involved with ID work, including its phases, challenges, and cognitive demands.

2.1 IDS Phases

Based on analysis from nine semi-structured interviews conducted with professionals who were responsible for ID work in their organizations, Goodall et al. [12] propose that ID can be broken into three distinct phases. The *monitoring* phase corresponds to the ongoing surveillance of an IDS, including sifting through the various alerts it generates. When monitoring reveals a potential security event, the *analysis* phase is initiated, which involves in-depth examination to determine if the alert is actually a security event. If a security event is confirmed, the *response* phase involves intervention and reporting of the event. Note that missing from this task analysis is IDS configuration. Thomson et al. [23] refine the Goodall analysis with data from two semi-structured interviews. They propose that, in addition to the above-mentioned three phases, ID work also involves a *pre-processing* phase. This phase occurs before the monitoring phase and corresponds to the actual IDS setup (e.g., configuring alerts, and/or generating filters for the alerts).

2.2 IDS Usability Challenges

Goodall et. al. [13, 12] propose that ID work is challenging due to expertise demands and its highly collaborative nature. ID requires significant expertise, both technical and organizational. Professionals need to have knowledge of their own unique network environment, since what is classified as a security event in one network may not be considered one in another network [12]. Attaining this degree of expertise is difficult, as much of the necessary knowledge is tacit and may be organization specific. Further complicating ID work is its collaborative nature that drives the need for practitioners to coordinate with other organizational stakeholders [13].

To obtain a fine-grained view of the challenges, Thomson et al. [35] use data from two interviews to perform a cognitive analysis of the three ID phases (pre-processing, monitoring, analysis, response). In general, they propose that all ID phases are challenging, but that the monitoring and analysis phases are the most cognitively demanding for practitioners. This high cognitive load derives from the need to integrate various sources of information in these two phases, including background knowledge on the network and the user base and information generated by the various tools involved in ID, such as the output of an IDS and network logs.

2.3 Support and Evaluation

IDSs generate large volumes of data, which subsequently security practitioners need to inspect. If this information is presented in textual form, as is the case for most of the existing commercial IDSs, then this places a high burden on the practitioners to make sense of the data. An alternative is to devise effective visual representation of the data to alleviate some of the cognitive burden and so facilitate the task of identifying security events (e.g., [22, 24]). For instance, the Intrusion Detection toolkit (IDtk) [22] generates glyph-based visualizations of network data, which may be raw packets or generated by an existing IDS, such as SNORT. IDtk uses color, spatial coordinates and glyph size to create the data visualizations, which aim to support the monitoring, analysis, and response phases of ID work.

To date, although studies have investigated the process of ID, very few usability evaluations of IDSs exist. One exception is Thomson et al. [36], who compare how different interface types (text vs. visual) support the monitoring and analysis phases through a laboratory experiment with 16 participants (2 professional ID analysts, 14 graduate students). The findings suggest that each interface type has its

respective strengths and weaknesses. For instance, a text interface provides access to fine-grained detail, affording flexible interactions and customizations; but it burdens the user with high quantities of data and the need to know the command syntax. A visual interface, on the other hand, can provide an overview of the data, which facilitates the detection of attacks; but it fails to provide fine grained detail and so some attacks may be missed.

3. METHODOLOGY

Prior work has shown the need for better security tools to detect malicious activity in networks and systems. These studies also propose the need for more usable tools that work in real contexts [20, 5]. To date, however, there has been little focus on the pre-processing steps of intrusion detection. We designed our study to fill this gap, as well as to further the understanding of IDS usability and utility, particularly as the IDS is installed and configured in an organization. Consequently, our research questions were:

- What do security practitioners expect from an IDS?
- What are the difficulties that security practitioners face when installing and configuring an IDS?
- How can the usability of an IDS be improved?

We used a qualitative approach to answer these questions, relying on empirical data from security practitioners who have experience with IDSs in real situations. Below we detail our data sources and analysis techniques.

3.1 Data collection

We collected data from two different sources. First, we conducted semi-structured interviews with security practitioners. Second, we used participatory observation, an ethnographic method [9], to both observe and work with two senior security specialists who wanted to implement an IDS in their organization. These two sources of data allowed us to triangulate our findings; the descriptions from interviewees about the usability of IDSs were complemented by the richer data from the participatory observation.

3.1.1 Semi-structured Interviews

For the HOT Admin project, we have conducted to date 34 *in situ* semi-structured interviews with 35 participants from various organizations (16 different organizations from 11 sectors, e.g., post-secondary educational, scientific services, financial services, consulting, manufacturing, insurance, and non-profit). All participants played a role in upholding security in their organizations; their positions ranged from IT manager to general IT staff to security staff. Each interview lasted approximately one hour. The interviews were subsequently transcribed and sanitized to preserve the participants' anonymity. During the interview, subjects were asked a variety of questions pertaining to the nature of security (e.g., challenges, tasks, tools, organizational influences, security culture, etc). Note that due to the diversity of participants' positions as well as the nature of semi-structured interviews, not all participants performed and/or discussed ID work. Information pertaining to the nine participants that did discuss ID is shown in Table 1.

Table 1: Participant Information (Semi-Structured Interviews)

ID	# Sector	Position
P2	Academic	Security Manager
P3	Financial Services	General Security
P4	Academic	General Security
P9	Academic	General Security
P12	Scientific Services	General IT
P15	Academic	General IT
P20	Academic	IT Manager
P23	Consultant	General Security
P24	Academic	General Security

3.1.2 Participatory Observation

The participatory observation was performed by the first author in one large, distributed post-secondary organization. It should be noted that the observer is a security specialist with four years of experience as a security consultant in a large telecommunications organization, although with no prior experience working directly with an IDS. To date, the observer has spent 15 hours working with two senior security practitioners who have worked together in the organization for several years, and are specialists in their areas, namely *servers* and *networks*. These two experts are in charge of the technical security projects in their areas, including the installation of an IDS. This project is currently at the stage where the IDS is connected to a production network, and is ready for tuning.

The participatory observation has consisted of two main activities: meetings and individual work. There have been a total of three, hour-long meetings between the two security specialists and the observer. The work on the IDS started with one meeting, followed by 12 hours of individual work, and continued with two further meetings. During the individual work, the observer had brief one-on-one interactions with the specialists to discuss specific issues related to IDS configuration. Throughout the process, the observer kept detailed notes of the meetings and interactions with the security specialists and of the IDS implementation.

3.2 Data analysis

The data from the interviews and participatory observation were analyzed using qualitative description [30] with constant comparison and inductive analysis. We first identified instances in the interviews when participants described IDSs in the context of the activities they had to perform. We next contrasted these descriptions with our analysis on the participatory observation notes. These notes were coded iteratively, starting with open coding and continuing with axial and theoretical coding [7]. Results were then organized by the challenges that the participants faced when deploying and maintaining an IDS system.

4. ANATOMY OF AN IDS

An IDS is a tool that detects abnormal behavior in systems. For the work reported in this paper, we are interested in those IDSs that monitor and detect attack patterns in network traffic. Such systems are commonly referred to as network IDSs. To monitor the networks, the IDS uses *sensors*, which are probes that are connected in the networks and that passively sniff the network traffic. To detect at-

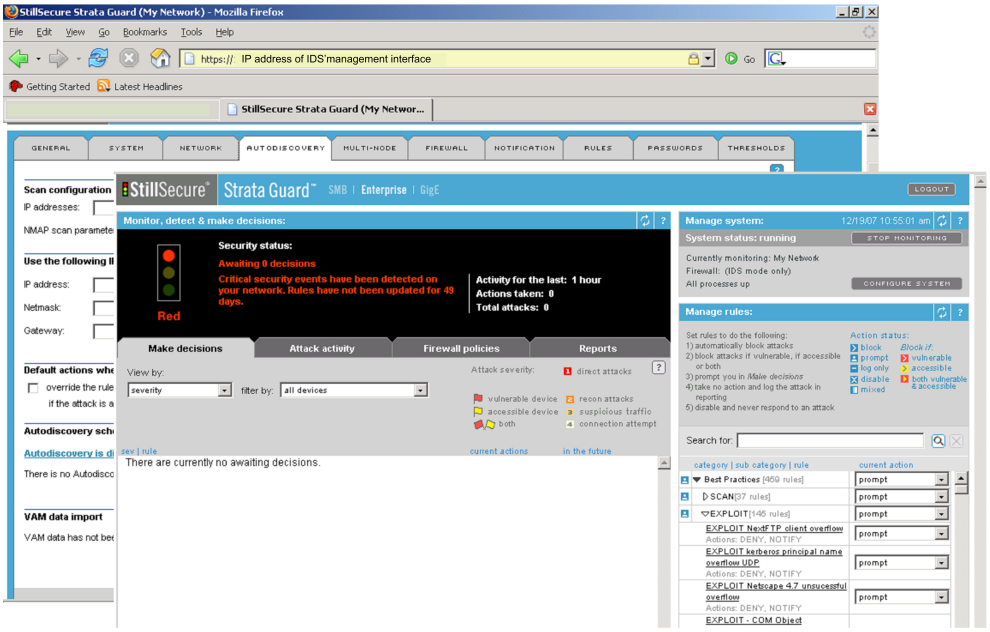

Figure 1: System configuration options of the IDS in the back. On top, Configuration options of the IDS's rules (bottom right) and status of alarms.

tacks, the IDS includes an *engine*, which typically performs detection via rules encoding attack patterns or signatures. Finally, the IDS provides mechanisms for administration, such as command line or graphical user interfaces.

4.1 The Deployed IDS

The IDS being deployed during the participatory observation was Strata Guard for small to medium businesses, version 4.5 [34]; the choice of system was based on a managerial financial decision. The IDS was acquired approximately five years ago. Since then, the organization has paid a maintenance to StillSecure (the vendor) for updates and general questions about the IDS's operation. Although current Strata Guard IDSs offer the option of being deployed with dedicated hardware (i.e., as an appliance), the version purchased by the organization came as a software package for general purpose servers. Another option, which was not available for the IDS version purchased, is IDS/IPS capability: (i) when operating as an IPS, the tool monitors and potentially intercepts network traffic (i.e., reacts instantaneously to attacks); (ii) when operating as an IDS, the tool monitors traffic and reporting alarms for off-line action.

The Strata Guard software included the following components: Linux operating system, PostgreSQL database, and a graphical user interface (GUI) as shown in figure 1, which enables the configuration of some but not all IDS settings (the IDS also includes a command line interface (CLI) that does enable practitioners to configure all aspects of the system). The support service provided by StillSecure gave immediate access to new attack signatures and also the option of opening trouble tickets in case of problems with the system.

During the participatory observation, the Strata Guard system was deployed as an IDS using software installed on an IBM server (Intel Xeon processor, 1 Giga RAM, 30 Giga Hard Drive). The server included two Ethernet ports: one used to monitor traffic, and one to manage the IDS server. To validate the IDS license and download rules to detect new attacks, the IDS needed to have access to the vendor's server (StillSecure) via the Internet, which was realized through its management Ethernet port.

5. INVESTIGATING IDS USABILITY

IDS usability evaluations should not be confined to the study of their graphical user interfaces: our data show that security practitioners also emphasize other factors (e.g., organizational) that influence the adoption of an IDS within an organization. We first highlight the main issues that security practitioners had to face during the integration of an IDS in a real network, as uncovered during the participatory observation. We then present the advantages and disadvantages of IDSs that participants described during the semi-structured interviews.

5.1 Issues Deploying an IDS

From discussions with the security specialists during the participatory observation, we learned that the initial objective for the IDS was to monitor traffic on the organization's internal networks. Alarms from the IDS were to be forwarded to the administrators of the appropriate networks. About two years prior to the participatory observation, the IDS had been installed by the security specialists in one particular network domain. However, it soon crashed, possibly due to memory space issues (the IDS GUI did not provide practitioners with functionality to manage the IDS's use of the hard-disk partitions), and/or from additional traffic from a newly-added wireless network. The former hypothesis related to memory issues was based on the fact that the default memory partition size was not large enough to accommodate the logs produced by the IDS; when a partition became full, it seemed the IDS started to overwrite other

system partitions not dedicated to the IDS. The security specialists did not have the time to confirm this hypothesis and analyze the exact cause of the system failure, so they decided to start again from scratch and install the IDS in another network. This re-installation was delayed for several months due to high workload and other priorities.

We next describe the main issues the security practitioners addressed and the decisions they made during the current IDS installation, which are distilled from the participatory observer's notes (see Appendix A for details). The issues include not only technical ones, but also human and organizational, providing a rich perspective on the challenges related to installing IDSs. As such, our findings may be useful for researchers and practitioners designing support for IDSs; they may also serve to guide the development of scenarios for evaluating IDSs in real contexts [29].

5.1.1 Deciding on the Purpose of an IDS

The target organization's main goal behind the adoption of the IDS was to complement the existing security controls (e.g., firewalls). The security specialists believed that the IDS would make monitoring of the organizational networks more efficient than other alternatives such as having to manually detect attacks via analysis of the firewall log files, using an IPS, or using an anomoly-based IDS. Manual analysis of firewall logs was deemed too complicated, time consuming, and had no guarantee of obtaining the consolidated attack reports the specialists needed. Automatically blocking traffic through an IPS was ruled out as it would have gone against the open culture fostered by the organization's academic nature. The specialists believed that an anomaly-based IDS would be less effective for their organization, as this organization involves a variety of security protocols and services, with highly irregular network traffic.

Monitoring malicious traffic was not the only purpose that security specialists had in mind for the IDS. They believed that the IDS could provide important statistics about the security of the network, and the security controls they had implemented in the network's boundary. Information about the number of attacks that actually crossed the organization's defenses could give the specialists not only a sense of the security of the internal systems, but would also provide support for proposing new security investments.

The purpose of the IDS was a critical factor influencing details of its deployment and use. For example, to test the security of the network's boundary, it would have been necessary to have a least two probes for monitoring the network, or two different IDSs located before and after the firewalls (see figure 2). However, the specialists did not know how to integrate the information from the two points, since it was not clear if the IDS provided functionality for doing so.

Given the limited resources available, the specialists decided to simplify the IDS installation as much as possible, and to install the IDS in the internal network only. We now describe their experience in doing so.

5.1.2 Constraints related to Integrating an IDS in the Network

Despite the fact that the security specialists had tried to simplify the deployment of the IDS by limiting its purpose, the IDS integration proved to be a challenging task, due to a number of organizational constraints. For example, to connect to the IDS, the specialists needed to have available

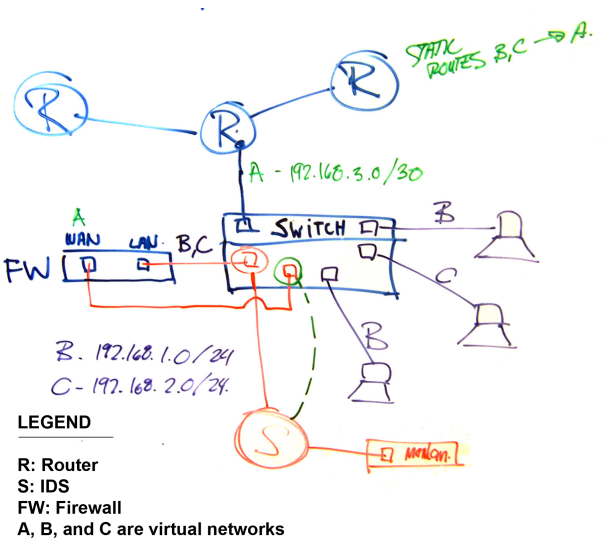

Figure 2: Network diagram used during one discussion about the installation of the IDS. The IDS has a connection to the management network and another to the port of the switch that transports internal traffic from the firewall. To compare configuration of the firewalls, it would be necessary to include another connection to the external traffic (dashed line).

ports (at least two in the case of the IDS used during participatory observation). In addition, they preferred to use the port mirroring feature of the switch connected to the IDS (see figure 2) to mirror traffic to the IDS, as this option provided the flexibility to select the traffic that they wanted monitored. These requirements became constraints for our participants, who could not find the necessary technical resources to connect the IDS in the critical network they wanted to monitor. Consequently, they decided to install the IDS into a less critical network; this decision was also influenced by other factors such as the distribution of IT responsibilities in the organization, as we explain in section 5.1.4.

5.1.3 Using A GUI for the Initial Configuration

Once the practitioners integrated the IDS into the network, the next step involved the installation of the IDS software. This required minimal intervention from the observer, who had to specify only the network settings and two passwords (one for the system and one for the internal IDS database). The GUI integrated with the IDS was intended to alleviate the burden of using a command line interface to administrate the IDS components (e.g., database, security engine) and to provide an easier method of tuning the rules. Specifically, the Strata Guard GUI provides an option (quick tune) to tune the system without the need of going rule by rule and considering the operating systems actually being monitored.

Although the participatory observer has not yet started the IDS tuning process, the initial configuration tasks have

revealed some of the shortcomings of the IDS's GUI. For example, the GUI does not allow the user to specify the hard-disk partitions assigned to the filesystem. This configuration option is important to the specialists, given that the pre-defined file space for the logs was too small when the IDS was used in the past. To manage log storage, an additional tool would be necessary. Similarly, the IDS does not provide support for configuring the IDS's security settings. Furthermore, the GUI does not allow users to configure the server's firewall rules, and so this task has to be done via the CLI, a task made difficult by the fact that the rules are non-intuitive and difficult to understand.

In general, although the GUI provided some support for configuring and maintaining the IDS (e.g., disable rules, take action on the alarms), the support was not adequate, given that the IDS was intended to work in a complex environment, influenced by the characteristics of the organization where it was going to be installed. The next section describes some of the key organizational factors influencing the deployment of the IDS.

5.1.4 Working Within a Distributed Environment

The observed organization was highly distributed in terms of IT administration, with various administrators in charge of different interconnected network domains. For these administrators, security usually was not the main priority. These two factors (distribution, security a low priority) triggered specific requirements that had to realized in order to integrate the IDS in the organization. For example, the monitored traffic flowed through various systems that were administrated by different practitioners. Notifications of the alarms the IDS detected in that traffic needed to be sent to the administrators of those systems, who should also be allowed to configure the IDS. Our participants hoped that the IDS would allow different levels of access depending on system characteristics, i.e., operating systems, IP addresses, specific network protocols. However, the deployed IDS did not provide such granularity to define access accounts.

Another issue related to distributed environment is the additional overhead it brings to the IDS project, which the security specialists wanted to minimize. The installation of the IDS in critical networks would have required the intervention of other specialists who administrated different sub-domains of those critical networks. These other specialists were not aware of the project from the beginning and might not have security as a first priority. This factor made our participants decide to discard the installation of the IDS in the critical networks. This decision resulted in a compromise, as the data may have been more interesting from the security point of view if these networks were included. This tradeoff between usability and utility is also discussed in the next section.

5.1.5 Balancing the Tradeoff between Usability and Utility

The security specialists required an IDS that was not only easy to use, but also gave relevant information about the security of the organization's systems. Consequently, the ideal situation would have been to install the IDS in the most critical network domain of the organization to generate meaningful reports about the security level of the networks, with a minimal use of resources. However, this did not occur; as discussed, organizational factors like distribution of IT responsibilities affected the decision to not involve critical networks due to the corresponding overhead of involving multiple administrators.

Another tradeoff between usability and utility was related to how the complexity of IDS configuration varied as a function of the network domains being considered for its installation. Specifically, the specialists could not tell how much more demanding it would be to install the IDS in a large network domain as compared to installing the IDS in a small network domain. This factor also affected the decision of where to install the IDS, as they believed that it would be much easier to install the IDS in a small network domain. However, it seemed that the only way to know how the complexity varied was to complete the full installation process on each of the candidate networks.

Another aspect that security specialists knew required a balance between usability and utility was related with the alarms the IDS generated. They knew that more false positives would require more time from them to investigate the alarms, thereby lowering the usability of the IDS. On the other hand, less false positives would imply less rules running in the IDS and, therefore, potentially more false negatives. Unfortunately, until the tuning process is complete and the IDS is in production, the actual tradeoffs between false positives and negatives will not be known.

5.2 Advantages and Disadvantages of IDSs

The results from the participatory observation have highlighted that there are more than just technical factors to consider when installing an IDS in an organization. In this section, we present our analysis of the interviews with various security practitioners, focusing on perceived advantages and disadvantages that IDS afford. As was the case with the results above, our findings span technical, human, and organizational dimensions.

As one of the participants from our field study stated, an IDS is *"one of the most controversial [tools]- some really love it, but some really hate it"* (P24). This controversy is likely rooted in the fact that IDSs have both strengths and weaknesses, and the tradeoff between the two is not always clear, as we discuss below.

5.2.1 Perceived Advantages

Our participants mentioned four key advantages of IDSs, including (1) problem identification, (2) monitoring with privacy, (3) decreased time pressure for maintenance, and 4) reduction of uncertainty.

The first perceived advantage is that an IDS can be a powerful tool to help identify problems (P4, P24). For instance, P24 stated that the IDS provided *"useful information about what kind of activities are outside a firewall and I want to have something inside the firewall too; to give me some idea whether something managed to go through"*. In identifying problems, an IDS *"makes good business value"* (P4).

Secondly, while security practitioners need to monitor their networks, they also need to maintain privacy of the organizational stakeholders. IDSs can support both of these goals. For example, one participant expressed how Argus [1] did so: *"Argus is a tremendous tool, it allows us to monitor activity and still respect privacy...because we're not looking at the data portions of the packets, on the header portions"* (P3).

Thirdly, security practitioners are notoriously overworked

and juggle a variety of tasks [5]. This sometimes means that they do not have the resources to attend to critical security tasks, such as ensuring that patching of systems happens in a timely manner. As a consequence, the systems become vulnerable and may even be compromised, something that occurred in one participant's organization. According to this participant, an IDS could help with this issue: *"we don't have to run around, for example tomorrow's... patch Tuesday. If we had this intrusion prevention we could patch quarterly. I don't have to run around and neither does anyone else"* (P14).

Finally, one issue that complicates security practitioners' work is related to the inherent uncertainty of their tasks. In particular, our participants mentioned that they are never certain as to the correctness of their activities (P3). An IDS could provide some assurances that everything is in order, e.g., *"...I am going to be considering keeping a closer eye on traffic both in and out, probably with an IDS, so that if there is something weird or not right going in and coming out, what have you, I can at least be alerted to it"* (P20).

5.2.2 Perceived Disadvantages

Despite the fact that an IDS affords advantages, some of our participants were hesitant as to its overall utility, which in turn discouraged them from adopting an IDS in their organization. The disadvantages that the participants mentioned included (1) the expense, (2) the degree of work and time required, (3) the unreliability of the IDS, and (4) the lack of clear utility.

The first disadvantage is that an IDS can be an expensive endeavor: *"so you can easily spend a quarter million dollars on an IDS and have 3 people running it"* (P4). This is exacerbated by the fact that security is often not a priority, and IDSs fall outside of the mainstream tools, i.e., *"[we do not have a commercial IDS because] we're tight budgetwise and security doesn't get a lot of budget outside of the main stuff, like anti-virus and firewall, and traffic shaper and stuff"* (P3).

Secondly, several of our participants stressed that IDSs are also costly as they require a lot of work and time resources (P3, P4, P9, P24, P12). This demand for resources happens both in the pre-processing IDS set-up phase and the monitoring and analysis phases. As far as configuration is concerned, tuning the IDS can be an arduous undertaking that requires both time and expertise: *"tools like Snort, they're great tools, but they require a lot of customization to get it down to something that understands your environment, so you have to turn alarms on and off based on what you're looking for, what's normal, what's not normal. When I first ran Snort in our environment I was getting thousands of flags a day"* (P9). A key issue with fine-tuning an IDS is to reduce the number of false positives (P4, P9, P24, P12), which occur when customization is not done properly. For example, one participant stated *"when I did run Snort in the past, which is looking for pattern matches on incoming traffic, it just had a ridiculous number of false positives"* (P12). Of course, fine tuning also means not blocking legitimate traffic (P3). Unfortunately, it is very difficult to determine how well an IDS is set up (P23). In the monitoring and analysis phases, lack of time was again an issue: *"I don't monitor that as much as I should be because of lack of resources, because it takes too much time... and then investigate the risks on [the IDS]"* (P3).

Thirdly, our participants sometimes found IDS software to be unreliable, which resulted in lost time and potentially important data, e.g., *"it's quite buggy and sometimes it would fill up all the log files so some partitions were filled up because of the humongous amount of logs ...it would just clog it up and you have to reinstall and then you can really kind of clean up the archive logs and stuff like that. It is just a nightmare"* (P24). Another participant mentioned that some IDSs sometimes dropped packets when they became overloaded (P2). This lack of reliability and potential for interfering with regular network traffic was a negative factor in participants' perceptions of the utility of an IDS.

Finally, although IDSs require many resources, their utility is not always clear. It is hard to see improvement in the security processes, *"you don't really notice any improvement"* (P4). Another consequence of the resources required to maintain an IDS is that often, they simply sit idle (*"we do have an intrusion prevention system in place but we haven't been using that effectively at all. It just kind of sits there and runs away"* P15).

6. DISCUSSION

Our findings suggest that the usability of an IDS is not solely determined by the usability of its GUI. We now discuss some of the associated human, organizational and technical challenges practitioners encounter when deploying an IDS, focusing on: (1) considerations before deploying the IDS; (2) the configuration and validation of the IDS; and (3) its ongoing usage. Where appropriate, we provide suggestions for addressing the challenges, which are based on three sources: participatory observation, interviews, and guidelines from the literature. While some of these challenges may not have obvious solutions, it is important that security practitioners, tool developers, and researchers are cognizant of the complexity of this process.

6.1 Considerations before deploying an IDS

There are number of challenges that impact an organization's decision to use an IDS. First, our interview analysis revealed that IDSs have not gained the same popularity as other *de facto* security tools, such as firewalls. This makes it more challenging for security practitioners to obtain management buy-in. This challenge could be alleviated with concrete data demonstrating an IDS's utility, however, obtaining the data is difficult for two key reasons. First, in order to obtain the data, the IDS needs to installed and configured within an organization, as generic reports may not reflect a given organization's characteristics. Second, once an IDS is installed and configured, the data needs to be transformed to a form readable by various stakeholders, including managers. To alleviate the latter challenge, an IDS should include reporting functionalities that tailor the information according to a user's specific needs. Furthermore, it should provide the ability to compare the outputs of different IDSs or IDS probes. This functionality would allow security practitioners to compare the state of security before and after the implementation of the IDS (a general version of this guideline is suggested in [26]).

Second, the decision to use an IDS impacts many stakeholders within the organization. These stakeholders need to be involved in the process, to maximize both the stakeholder buy-in as well as the benefits of installing such a tool. However, doing so comes with a cost due to the over-

head needed to manage the involved parties. Consequently, organizations may opt to reduce this overhead, even though this reduces the IDS utility (as was the case for the organization involved in participatory observation). Third, IDS configuration and use requires extensive resources from security practitioners, who typically have other competing priorities. Fourth, our participatory observation revealed that the installation of an IDS requires the participation of security specialists with knowledge and experience not only in network protocols and systems, but also about the organization itself. The observed security specialists had detailed knowledge of the organization, the networks that provided critical services for clients, and even clients' usage patterns.

The last three challenges derive from lack of security budget, tight schedules and security as a low priority in organizations [37]. To alleviate these challenges, one of our participants proposed that organizations planning to install an IDS should formalize the process via a dedicated project that includes allocation of resources and the responsibilities of the stakeholders involved: *"So we have internally a project approach...- it's going to have some people allocated to it and a certain amount of capital budget. Well then we write it up in a project and it goes through a project approval process through our senior management team."*(P15). Two other participants suggested allocating some dedicated and uninterrupted time for the IDS (P24, P9). To address budget issues, one participant proposed the use of open source tools (P19), an approach suggested by [25]. Such tools can afford benefits [28], such as better internal engines (P19, P25); however, our participants believed that these tools suffer from weaker reporting capabilities (P19, P22, P25) and less management buy in (P19), as compared to commercial tools.

6.2 Configuring and Validating an IDS

Once an organization makes the decision to use an IDS, the IDS needs to be installed and configured. Our participatory observation revealed a number of challenges related to these steps that we discuss below, along with guidelines to address them.

6.2.1 Collaboratively Evaluating Tradeoffs

One of the main challenges described by our participants during the IDS configuration process was the need for both broad and deep knowledge of services and organizational goals. Without this knowledge, it is difficult for practitioners to weigh the tradeoffs between increased ease of monitoring through a reduction in the number of false positives and the subsequent reduced IDS utility, due to increased false negative.

To obtain this knowledge, the installation of an IDS in the network requires collaboration with different experts in the organization. Our participatory observation showed cooperation between at least two experienced security specialists from the network and server areas respectively.

6.2.2 The Configuration Hurdle

Hill [17] states that the big hurdle for most users of security tools is not the user interface, but rather acquiring and installing the software. For the security specialists we observed, a factor complicating the IDS installation was uncertainty: they found it very difficult to predict the degree of effort that would be required to configure the IDS in a particular network. In the end, they found it necessary to go through the full installation process to determine the costs and benefits of the different configuration options according to the utility of the events the IDS detected and reported. This characteristic implies that an IDS might be classified as an "all or nothing" security tool, which makes its adoption and use in the organization difficult. This contrasts with other security tools that do not require intensive use of resources in their configuration to assess their benefits. For example, a security scanner can work with its default configurations and still generate useful reports on system vulnerabilities.

Since the configuration of an IDS is the breaking point for many potential users, IDS designers should aim to minimize the resources required to install and configure these tools. The Strata Guard system used during participatory observation provided several features in this direction, such as automatic discovery of the network's devices and a quick tuning option. However, its GUI did not allow the configuration of all the options required to optimize IDS usage (e.g., memory partitions). Furthermore, error messages the IDS generated during the installation were not helpful. Based on these observations and prior work, the following three guidelines aim to improve the usability of IDSs. First, IDSs should provide facilities for quick configuration, which can be realized, for instance, by grouping related parameter values [14]. Second, IDSs should provide meaningful help during the configuration process or ongoing usage [25]. Third, IDSs should provide documentation on the configuration process [14].

6.2.3 Determining an Appropriate Test Bed

A challenge our participants encountered during the installation and configuration process was determining an appropriate test bed environment for the IDS. In general, an IDS must be installed in a real environment to have a sense of its benefits; however, inserting the IDS into a production system might be difficult when there are other stakeholders involved who do not see the benefit of altering the networks.

To deal with the complexity of validating IDS configuration, one participant suggested first testing the IDS in a smaller network than the target one, so as to reduce the amount of traffic security practitioners has to contend with when testing: *"we have to redeploy it to a smaller network ... because it used to be on huge networks [and] we had tons and tons of traffic and tons and tons of ... alerts ... [it was] just too much"* (P24). This participant found that testing on a smaller network *"worked quite well"*, as it provided some useful information on network activities. What P24 suggested is a practice called "planning and rehearsal", as advocated in [2, 3].

If an IDS is installed in a rehearsal environment, the tuning will fit that network, but the tuned system may not fit the target environment. This issue highlights the complexity associated with IDS usage. More research is needed to better understand the trade-offs between smaller rehearsal environments to test an IDS, and the configuration impact of moving them to more complex networks that often transport the critical traffic in the organizations.

6.3 Ongoing Usage

After an IDS is installed and configured, challenges remain that impact its ongoing usage.

6.3.1 Monitoring an IDS

As discussed above, improving both the back-end of the IDS as well as the visualization of pertinent information for the practitioners monitoring the IDS alerts are active areas of research. In this vein, one of our participants explicitly discussed the need for improved recognition of anomalous network behavior via an IDS that had *"a bit of smarts"*, one that could watch and recognize trends over time (P3). This participant also described how without this ability, an IDS requires more human attention, as it generates alerts for innocuous network traffic that falls outside of the average throughout the year (e.g., in an academic institution before the term starts, there is very heavy traffic coming from web registration). Related work also provides some suggestions to improve monitoring. First, echoing the above-mentioned participant, Thomson et. al. suggest IDSs should provide automatic detection of malicious traffic behavior, realized for instance via pattern recognition techniques [36]. Second, IDSs should provide facilities for practitioners to fine-tune thresholds for generating alarms as well as facilities for suppressing alarms selectively [14].

6.3.2 A tool that fits the distributed nature of information security management

During our participatory observation, we found that different security practitioners needed to access the output of the IDS, but that doing so was complicated by the fact that these individuals were distributed across the organization. To address this challenge, related work has suggested that an administration tool should provide a shared view of the system state to its users [14, 2, 3]. Furthermore, Barret [3] suggests that tools with a shared view should provide proper authentication and authorizations, to ensure access is granted only to appropriate stakeholders. We recommend extending this concept by having the IDS tailor the view according to the needs of a given stakeholder.

Similarly, to facilitate monitoring and alerting, Haber [14] suggests that monitoring tools should provide alarm generation with a configurable destination. This feature enables an IDS to send its alarms through different channels (email, SMS, etc.) to different stakeholders distributed across the organization. In addition, McGann [25] suggests that providing reports in hypertext format would ease the distribution of reports to security practitioners across the organization. Beyond just providing the option of sending alarms to different stakeholders, we recommend that an IDS also provide features supporting on-line collaboration among these stakeholders. The IDS used during participatory observation could be configured to generate alarms using different communication channels (e.g., e-mail, SNMP), but it did not provide support for real-time collaboration (e.g., to discuss an alarm).

6.3.3 Reporting

We found reporting to be an important feature of an IDS. Reporting can demonstrate the economic value of the tool (not supported by the version of Strata Guard IDS in the participatory observation). It can also ease the burden of monitoring. For example, one participant described first deploying Snort to monitor the network. However, due to weaknesses of its reporting engine, his organization opted to acquire a commercial solution with better reporting features. The IDS should generate reports that help practitioners investigate the alarms. Furthermore, the IDS can help practitioners prioritize their tasks, by assigning priorities to alarms, or assigning each alarm to a practitioner for further investigation [37].

More flexible reporting has been recommended for security tools in general [5]. Flexibility can be afforded along a number of dimensions. As mentioned above, reports should be tailored according to the needs of the specific user reading them (e.g., manager, practitioner). Other options that may increase the utility and usability of reports include supporting a hypertext format [25] and using dynamic filters to help practitioners analyze large reports easily [10].

7. CONCLUSION

Intrusion detection systems are complex and provide many challenges for security practitioners. Prior IDS research has focused largely on improving the accuracy of these systems and on providing support to practitioners during the ongoing task of monitoring alerts and analyzing potential security incidents. One area that has received little attention is the pre-processing phase of IDS, but the installation and the initial configuration of an IDS can be so challenging that they can serve as a barrier to use. In this paper we have provided an investigation of these challenges through semi-structured interviews and participatory observation of one such deployment. Our analysis provided insights into the expectations that security practitioners have for an IDS, identified the difficulties they face when installing and configuring an IDS, and provided recommendations for improving the usability of ID systems.

One limitation of our work is that only 9 participants from the semi-structured interviews specifically discussed intrusion detection. Furthermore, two thirds of them came from academic organizations, as did those involved in the participatory observations. Although we argue that many of the issues around the deployment of IDS are organization independent, additional data from different organizational types would strengthen our results. Consequently, one aspect of our future work is to confirm and generalize the findings presented here. Additionally, we will begin to apply our findings towards the design of improved user interfaces for intrusion detection systems, focusing our attentions on relieving the burden on security practitioners that is inherent in configuring and maintaining an IDS. Until improvements are made across all phases of ID, it is clear that many security practitioners and organizations will continue to decide that the challenges of using an IDS will not be worth the effort required.

Acknowledgments

We thank the other members of the HOT Admin project for their feedback and our participants for taking part in our study. The HOT Admin project is supported by the NSERC Strategic Partnership Program.

8. REFERENCES

[1] Argus intrusion detection and prevention. http://www.qosient.com/argus/, February 2007.

[2] R. Barrett, E. Haber, E. Kandogan, P. Maglio, M. Prabaker, and L. Takayama. Field Studies of Computer System Sdministrators: Analysis of System Management Tools and Practices. In *Proc. of the*

Conference on Computer Supported Collaborative Work, pages 388–395, 2004.

[3] R. Barrett, P. P. Maglio, E. Kandogan, and J. Bailey. Usable autonomic computing systems: The system administratorsṕerspective. *Advanced Engineering Informatics*, 19(3):213–221, 2005.

[4] Base: Basic analysis and security engine. http://sourceforge.net/projects/secureideas, February 2008.

[5] D. Botta, R. Werlinger, A. Gagné, K. Beznosov, L. Iverson, S. Fels, and B. Fisher. Towards understanding IT security professionals and their tools. In *Proc. of ACM Symposium on Usable Privacy and Security (SOUPS)*, pages 100–111, Pittsburgh, Pennsylvania, July 18-20 2007.

[6] Bro intrusion detection system. http://bro-ids.org, February 2008.

[7] K. Charmaz. *Constructing Grounded Theory*. SAGE publications, 2006.

[8] S. Chebrolua, A. Abraham, and J. Thomas. Feature deduction and ensemble design of intrusion detection systems. *Computers and Security*, 24(4):295–307, 2005.

[9] D. M. Fetterman. *Ethnography: Step by Step*. Sage Publications Inc., 1998.

[10] S. Furnell and S. Bolakis. Helping us to help ourselves assessing administratorsúse of security analysis tools. *Network Security*, 2:7–12, February 2004.

[11] A. Gagné, K. Muldner, and K. Beznosov. Identifying differences between security and other IT professionals: a qualitative analysis. In *Proc. of Human Aspects of Information Security and Assurance (HAISA) (to appear, 10 pages)*, Plymouth, England, July 2008.

[12] J. Goodall, W. Lutters, and A. Komlodi. The work of intrusion detection: Rethinking the role of security analysts. In *Proc of the Americas Conference on Information Systems (AMCIS)*, pages 1421–1427, 2004.

[13] J. R. Goodall, W. G. Lutters, and A. Komlodi. I know my network: Collaboration and expertise in intrusion detection. In *Proc. of the ACM Conference on Computer-Supported Collaborative Work (CSCW)*, pages 342–345, November 2004.

[14] E. M. Haber and J. Bailey. Design Guidelines for System Administration: Tools Developed through Ethnographic Field Studies. In *Proc. of 2007 symposium on Computer human interaction for the management of information technology (CHIMIT), 9 pages*. ACM, 2007.

[15] K. Hawkey, D. Botta, R. Werlinger, K. Muldner, A. Gagne, and K. Beznosov. Human, organizational, and technological factors of it security. In *CHI'08 extended abstract on Human factors in computing systems*, pages 3639–3644, 2008.

[16] K. Hawkey, K. Muldner, and K. Beznosov. Searching for the Right Fit: Balancing IT Security Model Trade-offs. *Special Issue on Useful Computer Security, IEEE Internet Computing*, pages 30–38, 2008.

[17] A. Hill. Shortcuts, Habits, and Sand Castles. In *SOUPS '06: Proceedings of the second symposium on Usable privacy and security*, Pittsburgh, Pennsylvania, 2006. Invited talk.

[18] K. Hwang, M. Cai, Y. Chen, and M. Qin. Hybrid intrusion detection with weighted signature generation over anomalous internet episodes. *IEEE Transactions on Dependable and Secure Computing*, 4(1):41–55, 2007.

[19] K. Julisch and M. Darcier. Mining intrusion detection alarms for actionable knowledge. In *Proceedings of the 8th ACM International Conference on Knowledge Discovery and Data Mining*, pages 366–375, 2002.

[20] E. Kandogan and E. M. Haber. Security administration tools and practices. In *Security and Usability: Designing Secure Systems that People Can Use*, chapter 18, pages 357–378. O'Reilly Media, Inc., Sebastapol, 2005.

[21] G. Killcrece, K.-P. Kossakowski, R. Ruefle, and M. Zajicek. Incident management, 2005.

[22] A. Komlod, P. Rheingans, U. Ayachit, J. Goodall, and A. Joshi. A user-centered look at glyph-based security visualization. In *Proceedings of the IEEE Workshops on Visualization for Computer Security (VizSEC)*, pages 21–28, 2005.

[23] K. lynn Thomson and R. von Solms. Information security obedience: a definition. *Computers & Security*, 24(1):69–75, 2005.

[24] E. L. Malécot, M. Kohara, Y. Hori, and K. Sakurai. Interactively combining 2d and 3d visualization for network traffic monitoring. In *Proceedings of the 3rd international workshop on Visualization for computer security (VizSEC)*, pages 123–127, 2006.

[25] S. McGann and D. C. Sicker. An analysis of security threats and tools in sip-based voip systems. In *In 2nd Workshop on Securing Voice over IP*, June 2005.

[26] M. Nohlberg and J. Backstrom. User-centred security applied to the development of a management information system. *Information Management & Computer Security*, 15(5):372–381, 2007.

[27] Open source host-based intrusion detection system. www.ossec.net, February 2008.

[28] E. S. Raymond. The cathedral and the bazaar. *First Monday*, 3(3), 1998.

[29] J. Redish. Expanding usability testing to evaluate complex systems. *Journal of Usability Studies*, 2(3):102–111, 2007.

[30] M. Sandelowski. Whatever happened to qualitative description? *Research in Nursing & Health*, 23(4):334–340, 2000.

[31] K. Scarfone and P. Mell. Guide to intrusion detection and prevention systems (idps). Technical report, NIST: National Instutute of Standards and Technology, U.S. Department of Commerce, 2007.

[32] Squil. sguil.sourceforge.net, February 2008.

[33] Snort intrusion detection and prevention. http://www.snort.org/, February 2007.

[34] StillSecure. Strataguard ids/ips protection system. http://www.stillsecure.com/strataguard/index.php, February 2008.

[35] R. S. Thompson, E. Rantanen, and W. Yurcik. Network intrusion detection cognitive task analysis: Textual and visual tool usage and recommendations. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting (HFES)*, pages 669–673, 2006.

[36] R. S. Thompson, E. M. Rantanen, W. Yurcik, and B. P. Bailey. Command line or pretty lines?: comparing textual and visual interfaces for intrusion detection. In *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 1205–1214, New York, NY, USA, 2007. ACM.

[37] R. Werlinger, K. Hawkey, and K. Beznosov. Human, Organizational and Technological Challenges of Implementing IT Security in Organizations. In *Proc of. HASIA'08: Human Aspects of Information Security and Assurance (to appear, 10 pages)*, July 2008.

[38] R. Werlinger, K. Hawkey, and K. Beznosov. Security practitioners in context: their activities and interactions. In *CHI '08 extended abstracts on Human factors in computing systems*, pages 3789–3794, 2008.

APPENDIX

A. DETAILED NOTES FROM PARTICIPATORY OBSERVATION

A.1 Deciding on the Purpose of the IDS

During the first meeting with the two security specialists involved in the deployment of the IDS, one of the main discussion points was the type of reports that the IDS needed to provide. The server security specialist was looking for evidence to show the effectiveness of the rules that were implemented in the firewalls. To obtain this evidence, it was necessary to install two sensors in the IDS, one before the firewall and the other one after. The differences between the alarms shown by the two probes would give a sense of how well the firewalls were configured. Such a report would shed more light on the investment decisions and business cases that the organization was considering for IT security. For example, a report saying that no attacks were crossing the firewalls and routers would confirm that those devices were saving the organization money, by avoiding security incidents. On the other hand, if the firewalls and routers were not filtering properly, then this would provide our specialist with evidence to support the purchase of firewalls with better functionalities and centralized management.

The network security specialist was concerned about the IDS set up proposed by his colleague, the server security specialist, for two reasons. First, the IDS might be unable to process all the information from the probe set up before the firewall. Second, the information might have little use, as the priority is identification of attacks that could actually penetrate to the internal systems.

The final decision about the purpose of the IDS was determined by practical issues. Given the lack of resources (e.g., time, man hours), the IDS was going to be installed with its basic configuration, with one probe only. This decision was discussed in parallel with where to position the IDS in the network.

A.2 Integrating the IDS in the Network

The discussion in the first meeting described above was supported by a sketch on the whiteboard of the internal network, including the main routers, switches, firewalls, and servers (see figure 2). This diagram had two main objectives. The first objective was to reach a common understanding of the current status of the network. The importance of this shared understanding was evident during the discussion, as each specialist knew unique details about the network. The second objective was to find ports available for connection to the sensor and management IDS ports.

From a technical point of view, the decision about where to position the IDS had several constraints. One of them was the bandwidth of the critical traffic to be monitored, which had to be smaller than 100 Mbps. Another constraint was the routing necessary to reflect in one specific network-device port all the traffic to be monitored. To do so, the traffic had to go through different devices and links that may not have spare capacity. The decision about the location of the IDS was not made in the first meeting, and the discussion continued during the second and third meetings.

During the second meeting the connection of the IDS was discussed in more detail. The initial idea of connecting the IDS to one of the routers was deemed impractical as the network had been reconfigured with new devices. These new devices would require a special module (not installed at that moment) to mirror traffic in one of its ports. The other possibility was to connect the IDS with another device within the same network domain, but the only port available in that device for reflecting traffic was reserved for troubleshooting during the investigation of network anomalies. Within this option, there were also issues with physically carrying the traffic to the room where the IDS was going to be installed. The security specialists discussed if it would be possible to reflect traffic in one device in the middle, and then reflect again this traffic in a second device in the target room. This was deemed infeasible so they had to evaluate a physical extension of the cables to connect the IDS.

Several issues arose during the connection of the management port of the IDS. It was not clear if the IDS' management port should be in the management network or if it was necessary to create a different VLAN for it. The network specialist proposed to create another VLAN to connect the IDS, but this option was deemed too complex. Another issue was the security of the management port; in the case of a new VLAN, it would be necessary to configure additional firewalls specifically for the IDS.

Given the inconvenience of connecting the IDS's ports, the security specialists began to evaluate other alternative locations for the IDS. This change in location meant that they would be giving up monitoring the most important traffic in the network, but did have the benefit of decreased complexity. This situation would have an impact, as the IDS-related reports would include as interesting results as they were originally hoping. The final decision about the location of the IDS was postponed until the third meeting.

During the third meeting, the security specialists continued to discuss the option of installing the IDS in a less critical network. They finally decided to adopt this last option, connecting the IDS's sensor in the network that carried traffic generated by the organization's internal staff members. These conditions made the project less ambitious, and it was now considered a pilot study. The management port was connected to one production network. In making this decision, the specialists discarded the connection of this port in the management network, which carries all the management traffic from the organization's devices. The main reason for not taking this option was that the security specialists did not want to involve the administrators of the management network, in order to reduce the project overhead.

Another topic discussed in the meetings was related to the configuration of the IDS, described in the next section.

A.3 Initial configuration of the IDS

The security specialists knew from their previous experience that customizing the IDS to the connecting network is a time consuming and iterative process. An IDS that is well tuned should minimize both false positives (i.e.,alarms that correspond to valid traffic) and false negatives (failure to generate alarms for anomalous traffic).

The IDS configuration was done by the observer as part of his individual activities. To be more prepared for the eventual tuning in the real network, the objective for the observer's individual work was to become familiar with the IDS and its graphical interface. The first task was to reinstall the IDS software on the server. This process, which took 20 - 30 minutes, automatically installed the required components of the IDS: the Linux operating system, PostrgreSQL database, Snort rules for detecting malicious traffic, and the IDS graphical interface. The information required by the system to finish the initial configuration included: (1) the IDS port IP addresses, which were set for only the management port in one of the organization's internal, secure networks, and (2) two passwords, one for configuring the IDS and another for the database. The strength of these passwords was not checked by the system.

During the installation, the observer noted that the IDS did not allow for customizing the configuration of the system. This was not surprising, as this packaged IDS software is intended to alleviate the burden of having to integrate each of the IDS components manually, performing in the background all the necessary steps to have the system running quickly. However, there were some configuration options that the posterior use of the IDS showed needed customization. These options were related to: (1) the partitions that the system assigned to the filesystem, and (2) the server security settings that prevent unauthorized access of the IDS. These were not shown in the initial setup, and they were not accessible from the IDS's graphical user interface.

The partitions assigned to the filesystem were important because the security specialists knew from their previous experience that the file space for the logs might be too small. They wanted to check that the new version of the IDS had more space for the logs, but again this was a setting that could not be configured within the IDS' graphical user interface and that required the use of additional tools.

The ability to access the IDS security settings was important because the security specialists wanted to know what type of firewall rules, if any, were necessary to protect the IDS's ports. In its IPTable file, the IDS system recommended to not modify the default protection settings. These settings would be hard for a security practitioner to understand, particularly for one who was seeking the usability advantages afforded by the IDS's graphical interface.

Another drawback of the IDS installation and booting processes was that some error messages did not give sufficient information about their cause or consequences. For example, during the installation process, the message: "ACPI resource is not an IRQ entry" was displayed; and during the booting process, the message "smartd failed initialization" appeared. These messages became very relevant, as the IDS's management port could not initially connect to the central server of the vendor, and it was not clear this issue was due to problems related to those messages or to the configuration of the network's filters.

Troubleshooting of the IDS's Internet connection revealed that it was the network that was blocking the connections. As a consequence, the IDS's management port was moved to another, more open, network where the system started to download the rules from the vendor's server. The next step in the observer's individual work was to develop an understanding of the configuration options that the graphical interface provided, particularly the detection rules.

A.4 Effectiveness of the Graphical User Interface

Through the graphical user interface it was possible to make changes to the IDS rules and to the system's configuration (see figure 1). This last option allowed the modification of parameters such as the IP addresses of the ports, the autodiscovery option, and the networks to be monitored.

Without real traffic, it was very difficult to anticipate the types of alarms that the IDS was going to report. The only way to configure the system in such circumstances is by already possessing detailed knowledge about all the valid protocols that the network carried. However, the organization's open, distributed environment included traffic unknown to the security specialists. In such a situation, the organizational security policies may play an important role; for instance, a full set of rules could be disabled if the organizational policies do not exclude certain traffic (e.g., disable rules associated with port scanning).

This inability to anticipate alarms made it clear that the tuning process could not be done off-line; it was necessary to look at real traffic. Unfortunately, predicting the complexity of configuring the IDS in a particular network is very difficult. Consequently, the security specialists did not know if it was worth tuning the IDS for a simple domain of the network (i.e., low traffic, not many different types of devices) versus directly tuning the system for the more important, complex domains. The IDS interface also provided an option of quick tuning, which looked like a good way of avoiding the specification of all the default rules of the IDS (more than 1,000). However, without real traffic it was impossible to assess the tradeoffs associated with this option.

Another aspect important for the security specialists was the ability to notify other administrators about malicious traffic in their networks, as we now describe.

A.5 Configuring for Multiple Stakeholders

The IDS was supposed to detect security events and send alarms to those internal stakeholders who should be notified of security incidents. The security specialists were worried about the benefit of these notifications; they had to be very careful to limit the number of false positive notifications. This meant that the alarms issued by the IDS needed to be preprocessed.

Another functionality that security specialists needed in their collaborative environment was the definition of access accounts to the IDS with different privileges. For example, some users should be able to look at alarms from specific network domains, without looking at alarms from other domains. However, despite the fact that the IDS was monitoring traffic that was going to different domains, the system did not allow different accounts when it was installed with a single sensor node.