The Work of Intrusion Detection:
Rethinking the Role of Security Analysts

John R. Goodall
Department of Information Systems, UMBC
jgood@umbc.edu

Wayne G. Lutters
Department of Information Systems, UMBC
lutters@umbc.edu

Anita Komlodi
Department of Information Systems, UMBC
komlodi@umbc.edu

ABSTRACT

Intrusion detection (ID) systems have become increasingly accepted as an essential layer in the information security infrastructure. However, there has been little research into understanding the human component of ID work. Currently, security analysts face an increasing workload as their environments expand and attacks become more frequent. We conducted contextual interviews with security analysts to gain an understanding of the people and work of ID. Our findings reveal that organizational changes must be combined with improved technical tools for effective, long-term solutions to the difficulties of scaling ID work. We propose a three-phase task model in which tasks could be decoupled according to requisite expertise. In particular, monitoring tasks can be separated and staffed by less experienced ID analysts with corresponding tool support. Thus, security analysts will be better able to cope with increasing security threats in their expanding networks. Additionally, organizations will be afforded more flexibility in hiring and training new analysts.

Keywords

Information security, intrusion detection, field study, expertise, socio-technical systems.

INTRODUCTION

In this Internet era, organizational dependence on networked information technology and its underlying infrastructure has grown explosively. In conjunction with this growth, the frequency and severity of network-based attacks have also drastically increased (Allen, Christie, Fithen, McHugh, Pickel, and Stoner, 1999). At the same time, there is an inverse relationship between the decreasing expertise required to execute attacks and the increasing sophistication of those attacks – less skill is needed to do more damage (McHugh, 2001). Despite concerted efforts on preventative security measures, vulnerabilities remain. These are due to programming errors, design flaws in foundational protocols, and the "insider" abuse problem of legitimate users misusing their privileges (Lee, Stolfo, and Mok, 2000). Because of this, intrusion detection (ID), the sub-discipline of information security that monitors network events for signs of malicious or abnormal activity, has become an integral component in many organizations' approaches to security. While it may be theoretically possible to remove all security vulnerabilities through formal methods and better engineering practices, it is practically infeasible (Hofmeyr, Forrest, and Somayaji, 1998). Thus, ID will remain a crucial security practice for years to come.

Network intrusion detection systems (IDS) assist security analysts by automatically identify potential attacks from network activity and produce alerts describing the details of these intrusions. The type of network IDS that all of our participants use is signature-based. This operates by matching patterns of known intrusions or misuse against network activity to detect and classify attacks. The IDS is the last, and sometimes only, line of defense. If locks and safes in the physical world equate to preventative security in the virtual world, then an IDS is analogous to a burglar alarm. If the IDS produces an accurate alert, it affords the security analyst a last opportunity to respond before the damage is done.

The bulk of ID research has focused on improving the accuracy and coverage of IDSs. Surprisingly, there has been very little research into understanding and supporting the human analysts' ID-related tasks. The work of doing ID is an amalgam of complex tasks that requires extensive experience and knowledge. Analysts must continually monitor their IDS for signs of illicit activity. The sheer number of alerts generated by an IDS can be overwhelming. Current IDSs can trigger thousands of alarms per day, up to 99% of which are false positives (Julisch and Dacier, 2002). At the same time, analysts must keep up to date with changing configurations in their operating environment and newly discovered vulnerabilities or intrusion methods.

Although there are no in-depth workplace studies of security analysts, Yurcik, Barlow, and Rosendale (2003) do describe the complexities of system administrators' work based on the authors' first hand experiences at National Center for Supercomputing Applications (NCSA). This work frames the central ID challenge – the asymmetry between attackers and defenders, which gives the former the advantage in this escalating battle. Administrators must continuously identify and repair every vulnerability in their environment, while an attacker need only find a single vulnerable system to exploit. The attacker only has to hit once, while administrators must protect against all possible attacks.

Analysts responsible for defending their organization's network infrastructure face a difficult struggle to stay current with attackers' strategies and to be ever vigilant over their own environment. Their work is complex, requiring a great deal of expertise and experience in a particular environment. Finding individuals that fulfill these requirements is difficult. As networks grow and security threats increase, organizations will be hard pressed to find analysts that posses the requisite expertise to immediately accomplish effective ID.

Many researchers are working to provide support for these overburdened defenders, focusing on technological solutions to problems such as the excessive number of false positives and the inability of most IDSs to recognize novel attacks. These solutions, once available to practitioners in the field, will provide some temporary relief. To find long-term solutions, the underlying behaviors of ID work must be understood. Only then can appropriate "fit-to-task" technologies be designed and implemented. The results of this study offer a first step in this direction. Understanding the ID task decomposition, we have identified a viable socio-technical solution by rethinking existing staffing strategies and introducing a new class of IT support tools.

METHODOLOGY

Our research seeks to gain a situated understanding of the work of ID and the people who do it. To this end, we have conducted contextual interviews with information security experts to explore the mundane and exceptional processes of their day-to-day ID activity. The interviews were conducted *in situ* when possible, encouraging participants to demonstrate their interactions with their current IDSs and other support tools. Due to security policies, we were not allowed access to the actual working environment of all participants. In these three cases, the participants brought screenshots of their systems to demonstrate their interaction. The format of the interviews was semi-structured, generally following a prepared interview guide though allowing elaboration off topic. We conducted nine interviews that lasted from one to two and a half hours. Quotes from our participants are identified by number (P1-P9) in our text. The use of qualitative methods for this research was necessary to establish a rich description of ID work, notably the interactions between the human analysts and their systems. However, this approach forces a balance between population sample size and depth of data collection, in favor of the latter. All interviews and observations were transcribed and content coded, following grounded theory analytic protocols (Strauss and Corbin, 1998).

The participants had wide-ranging levels of ID experience, primary job duties, and organizational security needs. It was important to include a diverse cross-section of ID experts to ensure that our results were not limited to any single type of user or organization. All participants possessed a working knowledge of at least one IDS. Snort, an open-source, signature-based network IDS (Roesch, 1999), was the most common in a suite of often used tools. The participants' organizations and corresponding security practices ranged from the relatively open environments of university settings to the highly secure defense contractors, financial companies, and web service providers. The primary roles of the participants varied: most were network or systems administrators whose duties included ID, while two were dedicated information security analysts and two were IDS developers consulting for other IT departments in their organization. The participants' time spent on ID also varied greatly; for some, ID was a small part of their daily activities, but for others it consumed the majority of their time. Two security officers spent all of their time on security related tasks, of which ID played a major role; these were the only participants who had an inexperienced, small, but dedicated staff to assist in ID tasks. The two security researchers spent from half to three quarters of their time working on ID tasks, with a large chunk was devoted to new systems development and testing. The remainder of the participants, network or systems administrators, spent much less time doing ID as a part of their daily routine. For all, the amount of time for some tasks, particularly monitoring, was highly dependent on the environment and the nature of the IDS.

It is important to note that "most people aren't just analysts" (P3). Interacting with the IDS is just one part of a job that includes other systems, network, or security related tasks. This is particularly true in smaller companies where a dedicated security person is not likely to be cost effective because the organization does not believe their threat level to be high enough or the organizational security needs are limited. For example:

> That's how I describe myself now, more of a systems administrator who does security work, because my company isn't big enough to have a security person full time. Even with a hundred employees, I think I could spend my entire

day, every day of the year doing security stuff, making things better than they are now, but from the company's point of view, they don't need that. (P9)

This analyst works primary as the systems administrator for several dozen machines, but is also responsible for all of the organization's information security needs. This is typical of analysts being pulled in two directions by their organizational responsibilities.

DOING THE WORK OF INTRUSION DETECTION

Providing security requires an integration of tasks that include ID, preventative technologies for "hardening" systems, implementing encryption and authentication schemes, and educating users in safety-smart work practices. The work of ID involves more than reviewing IDS alerts and occasionally responding to critical events. ID itself cannot be accomplished effectively in isolation, it also requires monitoring and analyzing systems tangential to the IDS, as well as keeping abreast of the latest security information. Although there are subtle differences in how security analysts perform ID, all of our participants followed a similar process, which we analytically abstracted into three task phases: monitoring, analysis, and response. Each will be discussed in turn.

Monitoring

The first phase of ID work involves the ongoing surveillance of systems and network activity looking for indications of anomalous or malicious activity. This process is centered on the IDS, but is augmented by other monitoring tools and vulnerability scanners. In addition, analysts monitor an extensive set of resources, including web sites and mailing lists, for news of new attacks and vulnerabilities. These are the mundane daily tasks of ID. One analyst described how "keeping up with everything" constituted the majority of her time.

The number of alerts generated by most IDSs, especially on large, heterogeneous networks, can quickly become overwhelming. Because the role of the analyst usually involves more responsibilities than ID, coping with information overload forces a difficult choice. She can choose to limit the IDS signature set and thus the number of alerts, or be inundated with alerts to the point where she loses her ability to monitor the data on an individual level. An analyst who severely limits the signature set, can dramatically reduce the number of false positives. One analyst rationalizes this strategy as follows:

We also have a very limited signature set at the moment. Part of that is performance gear, the other part of that is just data inundation… If we were doing the alerts the way you should be doing alerts, which is, we don't have nearly as many things commented out [i.e., removed] as we do now. (P6)

However, limiting the IDS signature set is often undesirable because, although it may reduce the number of false positives, it is also likely to increase the number of false negatives, meaning that actual attacks go undetected. Those participants that did restrict their signature set did so with the knowledge that they were probably missing many actual attacks, but had no effective means of monitoring the large numbers of alerts. Pursuing the opposite strategy, analysts are forced to look only at aggregated summaries of alerts. In this case, choosing which alerts to pursue can devolve to almost random selection: "Generally I only pick one or two [alerts] of interest [to investigate]…based on what problems we've been having lately" (P1). Picking the one or two alerts of interest out of the hundreds or thousands generated is one way of dealing with the problem, but leaves many alerts unresearched.

Although the IDS is the primary means of detecting attacks, other monitoring systems and ad hoc data capture tools play an important role, from simple "pings" to determine if a server is listening to systematically collecting bandwidth and system usage statistics. These secondary monitoring tasks are useful in placing IDS alerts in a broader context, but also serve as a means of alerting analysts to "hotspots in traffic". For this, analysts create custom scripts to look at network data for "weird things", such as "the desktop machine that has ten times as much traffic as the next busiest machine on campus. You know there is something wrong there" (P1). This kind of custom designed system, a "very crude anomaly detector" (P1), is a useful supplement.

Simply keeping abreast of the latest attacks, vulnerabilities, and IDS signature updates could be a full time job in and of itself. All of the participants listed numerous web sites, mailing lists, and people that they regularly monitored for this kind of information. By far the most common resource was Internet mailing lists, which were continually watched for updated information that pertains to each analysts' idiosyncratic network environment, usually as part of the analyst's daily ritual: "my first stop every morning is the security websites to see what the threat du jour is and if there's something that we can craft a signature for if it's not [already available]" (P6).

Every analyst works in a unique environment, which is reflected in the diversity of their frequently read mailing lists. For example, participants in academia monitored a variety of lists specific to security in higher education, while the analyst who administered a number of Apache web servers on Linux machines monitored the mailing lists and web sites relevant to his versions of Apache and Linux. Mailing lists monitored by multiple participants were general incidence and vulnerability lists that attempt to quickly disseminate information about new bugs, vulnerabilities, and attacks. These high traffic lists, run by the community of information security experts, deal with a broad range of security issues for a wide variety of platforms and applications. They are general enough to pertain to nearly all environments, leaving it up to the analyst to determine which notifications pertain to their particular environment.

In summary, monitoring the status of the environment involves interaction with an IDS and other monitoring tools as well as following external information sources looking for vulnerabilities that match their particular environment. All of these monitoring tasks are part of routine ID work, time-consuming, but not as cognitively challenging as the subsequent analysis and response phases.

Analysis

The transition from the monitoring phase to the analysis phase begins with a security trigger event. For network monitoring, this event is usually an IDS alert or recognition of an anomalous event occurring in the environment, such as a sudden spike in traffic or user complaints of slow systems. Analysis of alerts involves not only the alert itself, but many sources of data that provide the contextual information necessary to determine whether or not the alert is an actual intrusion and if so, to assess its severity. For external resource monitoring (e.g., mailing lists), the announcement of a new vulnerability or attack method necessitates further research to determine its applicability and possible severity to one's network environment.

While monitoring is a part of the daily ID ritual, analysis and response are much more unpredictable, both in frequency and duration. Analysis could happen once a week, or, more likely, several times a day. If an IDS alert, network anomaly or new vulnerability is important, it could require the analyst to spend hours researching the problem before a diagnosis can be made. However, there are times when the experience of the analyst is leveraged to immediately dismiss an alert as a false positive:

> I know these criteria will always cause a false positive, even though there's different event types being triggered, you can always go and filter those out and you can just reassure yourself in two seconds that's another false positive. (P3)

> Certain IDS output…certain things you can always believe is a known attack, just because the experience you have and the rule signature may be so tuned to where it always detects that [attack]. (P8)

In the first case, the analyst can easily determine that an alert is nothing to worry about because of experience with similar alerts in his particular environment with this signature's criteria tells him that this alert is always a false positive. No further investigation is necessary. On the other hand, there are times when it is immediately apparent that an alert is not a false positive, based on his personal experience and intimate knowledge of the particular signature that generated the alert. If a quick glance does not reveal the importance of the alert his investigation continues, relying upon personal experience and his understanding of the alert's context.

Analysts' expertise includes general knowledge of network protocols and ID, but most importantly, knowledge of their unique network environment, because what is normal activity in one environment may be indicative of illicit activity in another. All of the participants echoed the importance of having an intimate knowledge of their particular context: "I can tell if something is going on, I know the network" (P3). This results in a very steep learning curve for ID analysts. One must not only learn the intricacies of network protocols and system operations, but how those are manifested in a particular environment that is constantly changing. Keeping up with changing configurations in the operating environment is difficult, but necessary to provide the context needed to analyze and diagnose an alert. This includes knowing the details of each machine on the network. This can be as simple as recognizing that a Windows IIS web server attack targeted at a Linux machine running Apache is clearly a false positive, though it becomes more challenging very quickly as the number and diversity of machines on a network increase. For the analysts we interviewed, tracking this context is accomplished through personal memory, without any external support. (Several participants did note that creating a database of this information would be helpful.) The following example demonstrates the importance of keeping track of the environment:

> An attack against an IIS-based web server is, that is a pretty severe alert, but if I have a filter that sits on that web server that just throws those things aside before it gets passed to the web engine, then I consider that a false positive. You know, me knowing my environment I can say, this web server, I set up this web server, it's fine. (P3)

When an alert describing an IIS web server attack is actually targeted at a host running an IIS web server, this should be of immediate concern. However, the analyst knows that the particular targeted machine has a filtering mechanism in place that cancels the particular request that the alert was generated on. The analyst knows this information because he set up the web server; if another analyst were to get the alert, he may have a very different reaction, since in and of itself it "is a pretty severe alert." For many of the participants, this detailed knowledge of the environment was manageable enough that they could do a reasonable job of using their memory to recognize certain target machines or services as being vulnerable or not to a particular attack. Obviously, as the size of the network increases and the systems and network administration tasks are more distributed, relying on personal memory for all of the details necessary for ID analysis becomes untenable.

The particular context of an alert is an important factor in IDS analysis. This contextual information is temporal; if the analyst does not have a mechanism for capturing the information quickly, it could be lost. The importance of "knowing the environment" is coupled with the problem that particular aspects of a network are always in flux. So, if the analyst knows, for example, that a particular machine is supposed to only be running a web server, but that other services are actually running, this is probably indicative of an attack. For example:

> What operating system is this [host] running, what services is it running? Not just what services is it supposed to be running, what is really running? (P7).

In order to gain this context, analysts rely on a myriad of data sources and tools that provide historical and current state information. These information stores must be accessed through separate tools and procedures, collated, and correlated back to the original data. Analysts often develop customized tools to facilitate this process. This process is challenging and often requires creative technical solutions, but necessary to form a complete and accurate description of the IDS alert or anomalous event. Performing the analysis of an IDS alert or vulnerability is grounded in the experience and expertise of the analyst, and in the relevant contextual facts surrounding the activity. Successfully diagnosing an alert or vulnerability is a difficult, complex task that requires an ability to improvise and develop custom tools and scripts to facilitate data collection and correlation.

Response

The most common forms of response in ID are intervention, feedback, and reporting. Intervention depends on the role of the analyst in the organization and organizational policies. Analysts who are also administrators of targeted machines would likely intervene themselves. The response to an attack in progress could be as drastic as unplugging a network connection: "probably the first thing would be unplug it" (P9). More common are responses that occur after the fact, such as patching the vulnerability or reinstalling the compromised machine from backup. Especially in larger organizations, the analyst in charge of ID is not the administrator of most machines. In this case, the response involves coordination among other administrators in the organization. Feedback is usually directed at the IDS or other elements of the security infrastructure. It includes tweaking or removing IDS signatures that generate an excessive amount of false positives, even if the signature was not guaranteed to always generate a false positive. As noted earlier, this practice is dangerous because it can lead to false negatives yielding undetected intrusions. Configuring and tweaking the IDS for the particular environment is one of the most challenging ID tasks, but one that teaches the analyst the nuances of that environment and how the IDS operates in that context. Feedback often involves submitting attack information to security mailing lists or vendors. Responses also include generating incidence reports for legal action and reports for management.

IMPROVING INTRUSION DETECTION

The problems of information overload and false positives are currently being addressed by a number of IDS researchers and vendors. Promising approaches include alert correlation (e.g., Debar and Wespi, 2001; Valdes and Skinner, 2001), data mining (e.g., Lee, et al., 2000), and alert clustering (e.g., Julisch, 2003). However, the problems are not only related to IDS. Recall that ID work involves multiple systems and information resources. True solutions must engage ID work holistically, involving all of its component parts. Technical solutions will help, but will not keep pace as the size of networks increase and security requirements become more stringent. Given our understanding of this as a socio-technical problem, this early research directs us to both areas of organizational change and tool development for an integrated solution.

Organizational Improvements: Decoupling ID Tasks

One organizational solution to the central problems of the ID analyst, the lack of time to analyze all alerts and vulnerabilities and the related information overload, is to decouple the core ID tasks. In particular, this would mean separating the monitoring tasks from the analysis and response tasks, then restaffing accordingly. An organization's security needs would be better served by taking advantage of an analyst's expertise for those tasks that actually require it and offloading the routine

work to those with less experience. Monitoring tasks are tedious, but are a necessary component of information security practice. However, in contrast with the expertise and contextualized knowledge required for the analysis and response tasks, the monitoring tasks are not nearly as cognitively demanding. Offloading the monitoring tasks would also allow less experienced analysts and administrators to become familiar with the nuances of their particular environment, in effect providing hands-on training in the area of ID that plays the central role in effective analysis: "knowing the environment." This apprenticeship would allow the novices to gradually come up to speed on the intricacies of a particular environment, without having to go through the typical "trial by fire" learning process that most of our participants described. By shifting the monitoring tasks away from the experts, their time would be freed up to focus on the challenges of analyzing, diagnosing, and responding to alerts and vulnerabilities, as well as to proactively "harden" their systems and networks. This would result in a more secure environment with the organizational resources shifted to the areas that best fit the requirements.

The junior staff tasked with monitoring could be trained off-site, learning the basics of networking protocols and ID. This would provide them with the essential knowledge of networking and security. Their subsequent on-the-job training would ground them in the particular details of their network environment. This contextual knowledge would be essential for them to mature from monitoring to analysis and response tasks. This task-based staffing would provide organizations with more flexibility in hiring new analysts, since they would not need to find analysts who have current ID experience in other organizations, experience that would not necessarily be immediately convertible to a new environment anyway.

Technical Improvements: Monitoring Support Tools

As described earlier, monitoring the IDS and secondary systems is essentially an anomaly detection problem. Similarly, searching external sources for relevant security related information involves matching the unique characteristics of one's network with the details of new vulnerabilities or attacks. These are tasks that an inexperienced analyst can be trained to do, particularly if given the right IT support tools. For monitoring the environment, information visualization can provide effective support due to its ability to highlight patterns and anomalies in large amounts of data. Information visualization is: "the use of computer-supported, interactive, visual representations of abstract data to amplify cognition" (Card, Mackinlay, and Shneiderman, 1999). Visualization takes advantage of strong human pattern recognition skills with visual representations that can make patterns and anomalies obvious to the user (Card, et al., 1999). One of our participants, who works at a university with several students available, describes the utility of such a visualization on a projection screen at the help desk:

> [Someone could be] glancing up at it occasionally, notice anything that appears to be funky, even if they don't call me, even if they just write it down in a log and say 'funky thing.' (P1)

In offloading the monitoring tasks, visualization tools need to be designed so that anomalies can be easily flagged for later analysis by more experienced analysts. Several researchers are working on applying information visualization principles to ID monitoring (e.g., Erbacher, 2002). Decoupling the tasks of ID will require new technologies to be developed to support less experienced analysts in accomplishing the monitoring tasks, some of which are already being prototyped. By combining organizational change with technological evolution, security analysts will be better able to cope with the increasing security threats in their expanding networks.

CONCLUSION

The work of intrusion detection is complex and challenging. Many researchers are searching for technological solutions to alleviate some of the problems in ID. However, these solutions may work best in conjunction with organizational changes. As the size and scope of organizational security requirements expand without a corresponding increase in staffing, an increased workload is placed on the security analysts. We identified three distinct phases of ID work: monitoring, analysis, and response. Because the analysis and response phases of ID are highly dependent on the expertise of the analysts, the monitoring phase lends itself to being offloaded to less experienced staff. This workforce can be trained to do the specific tasks associated with monitoring, interacting with new tools designed specifically to facilitate monitoring activities. This task separation, coupled with emerging technologies such as information visualization, applied to the particular problems of the monitoring tasks, can foster expertise in novices while giving the experts the time and resources to focus on those tasks that require their particular skills. We are currently designing information visualization tools to support the network monitoring and other ID tasks (for details, see: Komlodi, Goodall, and Lutters, 2004). We also plan to investigate the applicability of other technologies to fit the particular needs of each of the ID tasks.

REFERENCES

1. Allen, J., Christie, A., Fithen, W., McHugh, J., Pickel, J. and Stoner, E. (1999) State of the Practice of Intrusion Detection Technologies, *CMU/SEI-99-TR-028*.

2. Card, S. K., Mackinlay, J. D. and Shneiderman, B. (1999) *Information Visualization: Using Vision to Think,* Morgan Kaufman Publishers, San Francisco, CA.

3. Debar, H. and Wespi, A. (2001) Aggregation and Correlation of Intrusion-Detection Alerts, *Recent Advances in Intrusion Detection (RAID)*, 85-103.

4. Erbacher, R. F., Walker, K. L. and Frincke, D. A. (2002) Intrusion and Misuse Detection in Large-Scale Systems, *IEEE Computer Graphics and Applications,* 22, 1, 38-48.

5. Hofmeyr, S. A., Forrest, S. and Somayaji, A. (1998) Intrusion Detection Using Sequences of System Calls, *Journal of Computer Security,* 6, 3, 151-180.

6. Julisch, K. and Dacier, M. (2002) Mining Intrusion Detection Alarms for Actionable Knowledge, *Proceedings of ACM Conference on Knowledge Discovery and Data Mining*, 366-375.

7. Julisch, K. (2003) Clustering Intrusion Detection Alarms to Support Root Cause Analysis, *ACM Transactions on Information and System Security,* 6, 4, 443-471.

8. Komlodi, A., Goodall, J. R. and Lutters, W. G. (2004) An Information Visualization Framework for Intrusion Detection. *Proceedings of ACM Conference on Human Factors in Computing Systems (CHI)*, in press.

9. Lee, W., Stolfo, S. J. and Mok, K. W. (2000) Adaptive Intrusion Detection: A Data Mining Approach, *Artificial Intelligence Review,* 14, 6, 533-567.

10. McHugh, J. (2001) Intrusion and Intrusion Detection, *International Journal of Information Security,* 1, 1, 14-35.

11. Roesch, M. (1999) Snort - Lightweight Intrusion Detection for Networks, *Proceedings of Thirteenth Systems Administration Conference (LISA)*, 229-238.

12. Strauss, A. and Corbin, J. (1998) *Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory*, 2nd Edition. Sage Publications, Thousand Oaks, CA.

13. Valdes, A. and Skinner, K. (2001) Probabilistic Alert Correlation, *Recent Advances in Intrusion Detection (RAID)*, 54-68.

14. Yurcik, W., Barlow, J. and Rosendale, J. (2003) Maintaining Perspective on Who Is the Enemy in the Security Systems Administration of Computer Networks, *ACM CHI Workshop on System Administrators Are Users, Too*.