

Those Who Shield Others Are Users Too!

Experience from User Studies of Security SysAdmins

William Yurcik
National Center for Supercomputing
Applications (NCSA)

Urbana, IL
byurcik@ncsa.uiuc.edu

Ramona Su Thompson
University of Illinois at Urbana-
Champaign (UIUC)
Human Factors Division
Savoy, IL
ramonasu@uiuc.edu

Esa Rantanen
University of Illinois Urbana-
Champaign (UIUC)
Human Factors Division
Savoy, IL
rantanen@uiuc.edu

ABSTRACT

System administrators are users too and there have been few user studies of them in the security domain. This paper describes the challenges we faced in designing and performing studies of security system administrators using visualization tools as well as our experiences in addressing these challenges.

Categories and Subject Descriptors

C.2.0 [Computer-Communication Networks]: General – *security and protection*; C.2.3 [Computer-Communication Networks]: Network Operations – *network monitoring*; H.5.2 [Information Interfaces and Presentation]: User Interfaces – *graphical user interfaces (GUI)*; I.3.6 [Computer Graphics]: Methodology and Techniques – *interaction techniques*; K.6.5 [Management of Computing and Information Systems]: Security and Protection.

General Terms

Measurement, Design, Experimentation, Security, Human Factors.

Keywords

user study, usability, security, system administration, task analysis, intrusion detection, Internet security visualization

1. INTRODUCTION

Security system administrators have the unique role of maintaining the overall integrity of functional tasks within an organization. Without the proactive and quick reactive protection they provide, an organization would likely be subjected to numerous problems significantly decreasing productivity.

While there is a growing body of user studies focused on end user security tools usability, there are few user studies focused on security system administrators [2]. Security administrators, however, are users too, and they are the focus of our studies. We ar-

gue that since security administrators typically act on the behalf of many end users, improving the usability of their tools and procedures through user studies promises to have a larger impact on overall organizational security than end user studies alone. System administrators not only shield end users from larger problems but also decrease the need for end users to manage (or even know arcane technical details about) security on infrastructure systems they depend upon.

The security administrator's job of protecting an organization from attacks is difficult because it is high-volume, multi-dimensional, and dynamic, requiring high levels of expertise and the integration of numerous resources [3,4,5,19,20]. There is a consensus that the volume of information, repetitiveness of the tasks, and response times involved in these tasks beg the use of some type of automation—possibly even the replacement of the human-in-the-loop with automated agents. However, to date no automated system has yet approached the effectiveness of human security administrators.

While Internet security administration is unique due to its problem domain, there are other domains with similar challenges. A meta-analytic study of the human network monitoring for water systems, electric power grids, air traffic control, and nuclear power plants [11] show that visualization tools are a commonly used form of automation to support human operators monitoring activity on a network of systems.

Multiple visualization tools have been developed to aid Internet security system administrators [1,8,9,16,18,20]. However, there have been few user studies to quantitatively evaluate their effectiveness. No definitive user studies have shown that visualization improves effectiveness in this problem domain, in what scenarios it would be most useful, and how much improvement might be expected over traditional tools and interfaces.

In this paper, we share our experience in user studies of security system administrators aimed at finding the answers to these questions. We designed and completed user studies with security administrators at multiple levels including user requirements elicitation [20], task analysis [13], and performance measurement [12,14]. From our experience we present some views of the challenges of planning and executing security administration user studies and, where possible, insights about how these challenges may be overcome.

The remainder of this paper is organized as follows: Section 2 provides the context for our user studies and a brief background on security sysadmins tasks. Section 3 presents the challenges we

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ACM CHI'07 Workshop on Security User Studies: Methodologies and Best Practices, May 2007, San Jose CA USA.

Copyright 2007 ACM 1-58113-000-0/00/0004...\$5.00.

found in designing and performing our user studies as well as our experiences addressing these challenges. We end with a summary and conclusions in Section 4.

2. BACKGROUND

End users are dependent upon the availability of networked services such as authentication servers, storage systems, and transaction servers. The complexity of managing these services has grown tremendously with typically hundreds or thousands of components requiring skilled human effort to install, configure, upgrade, monitor, and debug. Some of system administration is automated; however, the actual degree of automation is much lower than many people assume [2]. Human operators are still very much in the loop, especially during emergencies.

Research has been conducted in understanding security system administration, specifically in the task of intrusion detection. Studies range from ethnographic studies [6] to requirements elicitation [20] to cognitive task analyses [3,4,5]. This research provides insight into the task of intrusion detection and a means for finding common themes.

Better system administration tools can produce substantial benefits with breakthroughs in five different dimensions: (1) tools to manage the scale of networked systems in terms of size and complexity, (2) tools to manage diverse networked systems independent of vendor-specific implementations, (3) networked system monitoring tools tuned to operator attention capabilities, (4) networked system troubleshooting tools leveraging human intuition, and (5) collaboration tools to aid human interaction especially during emergencies.

Toward support of security sysadmins along these five dimensions, visualization tools have been developed to reduce the cognitive workload of gathering and integrating information by the engineers [1,8,9,16,18,20]. These tools are typically very specialized; for example, they may require specific input data or only detect specific events. The latest generation of visual data mining tools and animated GUIs take advantage of human perceptual skills to produce substantial results, empowering users to perceive important patterns in systems, identifying components that need further scrutiny, and enabling sophisticated decision-making. But seeing information is only a start. Users also need to explore and manipulate information using real-time tools to zoom in and out of data, filter the data, and relate the information to other data sources, as well as having the ability to *undo* actions if they make a mistake. However, little empirical research has been conducted on the use of Internet security visualization tools and their effectiveness for the task of intrusion detection.

3. CHALLENGES

In designing and performing user studies in the context of security sysadmins, we found several challenges described in this section. The intention is to bring these challenges to the attention of other researchers so we can learn from each other.

3.1 Multiple Roles

One of our first steps is to identify and observe security sysadmins to learn what would be the most valuable research problems to study. We learned that the role of security sysadmin is not always clearly defined. Each organization is different and an organization may not have dedicated staff to monitor system security. It is ex-

actly when systems security staff balance multiple roles that system security may be put at greater risk because of divided attention. In order to overcome this challenge, research needs to be conducted to reveal exactly how people balance multiple roles across tasks and time. This can be done by examining the different staff roles, the different tasks to be juggled, how much time should be devoted to system security versus how much time is actually being devoted, what system security tasks are typically completed versus left uncompleted, etc.

3.2 Task Analysis

A task analysis of a security sysadmin should identify and qualitatively describe all tasks and subtasks performed (including those performed in parallel) along with their relationship to overall operational goals. While cognitive task analysis has shown that the task of intrusion detection can be divided into multiple subtasks [13], the extent to which each is performed varies across different organizations; each organization has their resources and processes tailored for their unique environment. This makes it difficult to (1) generalize about task analyses between organizations and (2) for a security sysadmin to learn or collaborate with another sysadmin. We suggest that more ethnographic studies need to be performed (such as [6]) in order to identify common tasks, resources, and processes across different organizations.

Based on our experience, we highlight the following common tasks which we feel are most valuable to study via quantitative measurements: (1) searching for known attacks with signatures, (2) discovering new attacks that have not yet been characterized, (3) situational awareness of systems so events can be placed in context, (4) the continuous integration of new tools to enhance search for different types of events, and (5) inter- and intra- organization collaboration between security sysadmins.

With numerous tasks and subtasks present at any time, we found that it can be difficult to pinpoint the end of one task/subtask and the beginning of another, especially when making precise empirical measurements. Analyzing our results, we did find a clear demarcation of tasks in which humans perform better than machines (discovery versus tasks in which machines perform better than humans (search) [7]. While there are few tasks that fall exactly at the polar ends of this spectrum (human-only with no automation at one end of the spectrum to automation with no human-in-the-loop at the other end of the spectrum), we find this insight a powerful guide for designing tools to leverage both human and machine capabilities.

3.3 Intervening Variables

There are many intervening variables to account for when designing and performing security sysadmin user studies. The following three are typical for most user studies: (1) user aptitude for different tasks, (2) user experience on different tasks, and (3) user training.

In our experience we did find users with varying skill and familiarity with the tasks being tested [14]. We measured the difference in results between the different groups with the goal being no significant difference. We address the issue of training by providing users with expert knowledge through cheat sheets and answering questions during the testing (which are recorded).

We found that the following intervening variables were more specific to our studies of security sysadmins and visualization tools:

- (1) Artifacts of specific implementation of the visualization tools tested (strengths/weaknesses, well-designed, easy to use);
- (2) Test cases and scenarios used in the study (biased for/against different tools, artificial vs. realistic); and
- (3) User expertise.

User testing for visualization of Internet security requires specific implementation(s) to be tested. Each implementation has different design characteristics such that results from testing one implementation may not hold for other implementations of the same functionality. When a specific implementation is tested, the strengths and weaknesses of that specific tool should be disclosed to put results in context. Ideally, multiple implementations should be tested so consistent results can be identified but this may not be practical due to finite development and testing resources. Testing specific implementations does provide valuable results that can be compared in later studies.

Independent of different implementations, the selection of specific test cases or scenarios can bias user study results. The goal is to select attack test cases that are fair, that is, representative of attacks that are seen across different organizations. For example, Internet security visualization tools typically help discover scanning activity (due to geometric graphical patterns) and denial-of-service attacks (due to significant changes in depicted traffic volumes). However, Internet security visualization tools typically do not highlight masquerade attacks, spoofing attacks, non-scanning worm/virus propagation or phishing attacks. Therefore, care should be taken to uniformly distribute test cases across a problem domain or test case bias should be disclosed to put results in context. One technique we have used is to distribute test cases based on rankings from independent authorities (e.g., SANS¹, CSI², and the HoneyNet Project³).

Given a fair selection of test cases, the data used in each test case should be realistic, even when real data cannot be used. In our work, we have taken traffic data captured from real attacks and then used anonymization tools to both obscure sensitive information and limit the search space while still providing enough information for testing within a reduced scope.

Perhaps the hardest challenge for user studies of sysadmins is varying expertise of study participants. Security sysadmins tend to gain their expertise from on-the-job-training, and the speed at which they gain expertise varies between individuals. This challenge is further exacerbated by the fact that there are little formal training conducted by organizations (or by the professional community) since environments tend to be specialized for general training.

¹ SANS (SysAdmin, Audit, Network, Security) Institute is a training organization as well as a clearinghouse for timely Internet attack information <<http://www.sans.org/>>.

² CSI (Computer Security Institute) is a training organization that also conducts an annual survey of computer crime in cooperation with the San Francisco Federal Bureau of Investigation's (FBI) Computer Intrusion Squad <<http://www.gocsi.com/>>.

³ The HoneyNet Project is a non-profit research organization that shares Internet attack information found within its network of honeypots (as well as sharing tools and techniques) <<http://www.honeynet.org/>>.

3.4 Evaluation Methods

Are controlled laboratory user studies the best way to understand sysadmins using Internet security visualization tools? Laboratory user studies provide quantitative results that can be replicated and the results generalized. However, laboratory testing removes the situated nature of most sysadmin work activity where interruptions, support material, and social collaboration are the norm [10]. Most studies of security sysadmin have involved less than a dozen users of which a majority may be students and not sysadmins.⁴ Ideally a user study should test tasks representative of what real security sysadmins actually perform across different organizations. Adapting tasks to a controlled laboratory setting with small number of students as surrogates for sysadmins may not be the best way to capture this.

Shneiderman and Plaisant [10] advocate the use of more Multi-Dimensional In-Depth Long-term Case studies (MILCs) as better suited to study the creative activities of users of information visualization tools engaging their own problems. So, as an alternative, sysadmin use of Internet security visualization tools may be evaluated using MILC based on usage (observations, interviews, surveys, automated logging) over 1-3 years and expert user success in achieving professional goals [10]. As a practical constraint, MILC studies are clearly a labor-intensive alternative requiring larger research teams to complete.

Just as studying security sysadmins requires adaptation in research design, so does studying visualization. Information visualization is a high-level cognitive activity that involves users testing hypotheses by looking for patterns over time. Users typically need to visualize the same data from different perspectives with different tools collaborating with others over a long period of time [10]. Important discoveries may be rare but when they do occur they can provide important new insights, possibly even shaking common beliefs [10]. Studying creative visualization processes using MILC techniques is appealing since controlled laboratory and tool experiments may appear too limiting in comparison. MILC and related ethnographic studies do have weaknesses in that observations may be misinterpreted, important events overlooked, and normal behavior changed under observation – but these same problems are also found in laboratory-based studies.

It is our judgment that a combination of laboratory experiments and MILC techniques are needed to best capture the impact of visualization tools with security sysadmins. For example, within weeks an iterative process can be accomplished including quantitative laboratory measurements of important tasks, replication by other researchers, and processes refined. While MILC takes significantly longer, its focus on benefiting users in their natural setting makes it a set of research methods we intend to pursue in future work. For instance, despite evidence supporting security visualization tools over security command line tools we found factors in the natural sysadmin environment which changed our integration plans [17].

4. CONCLUSION

Security sysadmins protect other users and thus are important to study due to their disproportionate impact on an organization's Internet security. From our experience, we have highlighted dif-

⁴ The seminal study of PGP usability tested only twelve users [15].

ferent challenges in designing and performing studies of security sysadmins using visualization tools. While ‘the devil is in the details’, investment to improve HCI studies in Internet security promise to enhance our understanding of how best for humans to manage security in the different roles they serve (end user, sysadmin, etc). We encourage others to build on our experience.

5. ACKNOWLEDGMENTS

First, we would like to thank the many participants in our user studies (predominantly UIUC computer science graduate students). Many of our ideas follow directly from their observations. Second, we would like to thank the primary software developers of the Internet security visualization tools we tested, Xiaoxin Yin/UIUC for *VisFlowConnect-IP* and Ratna Bearavolu/NCSA for *NVisionIP* – their patience was extraordinary when negotiating change modifications based on user feedback. Third, we would like to acknowledge Ben Shneiderman and Catherine Plaisant (both University of Maryland) for their input about the tradeoffs between case studies versus controlled experiments.

6. REFERENCES

- [1] Abdullah, K. et al., IDS Rainstorm: Visualizing IDS Alarms, *IEEE VizSEC*, 2005.
- [2] Barrett, Y-Y. M. Chen, and P.P. Maglio, "System Administrators are Users, Too: Designing Workspaces for Managing Internet-Scale Systems," *ACM CHI Workshop: System Administrators are Users, Too*, 2003.
- [3] D'Amico, A., and M. Kocka, Information Assurance Visualizations for Specific Stages of Situational Awareness and Intended Uses: Lessons Learned, *VizSEC*, 2005.
- [4] Goodall, J.R., W.G. Lutters, and A. Komlodi, I Know My Network: Collaboration and Expertise in Intrusion Detection, *ACM Conference on Computer Supported Cooperative Work (CSCW)*, 2004, 342-345.
- [5] Goodall, J.R., W.G. Lutters, and A. Komlodi, The Work of Intrusion Detection: Rethinking the Role of Security Analysts, *Americas' Conference on Information Systems (AMCIS)*, 2004, 1421-1427.
- [6] Kandogan, E. and E. Haber. *Security Administration Tools and Practices*. In Cranor, L. and Garfinkel, S. (eds.) *Security and Usability: Designing Secure Systems That People Can Use*, O'Reilly, 2005, 357-376.
- [7] Lakkaraju, K. et al., Closing the Loop in NVisionIP: Integrating Discovery and Search in Security Visualizations, *IEEE VizSEC*, 2005.
- [8] Lakkaraju, K. et al., NVisionIP: NetFlow Visualizations of System State for Security Situational Awareness, *ACM VizSEC/DMSEC*, 2004.
- [9] McPherson, J. et al., PortVis: A Tool for Port-Based Detection of Security Events, *ACM VizSEC/DMSEC*, 2004.
- [10] Shneiderman, B, and C. Plaisant, Strategies for Evaluating Visualization Tools: Multi-dimensional In-depth Long-term Case Studies, *BELIV Workshop of the ACM Advanced Visual Interfaces (AVI) Conference*, 2006.
- [11] Su. R., and W. Yurcik, A Survey and Comparison of Human Monitoring of Complex Networks, *10th Intl. Command and Control Research and Technology Symposium*, 2005.
- [12] Thompson, R. S., E. M. Rantanen, W. Yurcik, and B. Bailey, Command Line or Pretty Lines? Comparing Textual and Visual Interfaces for Intrusion Detection, *ACM Conference on Computer/Human Interactions (CHI)*, 2007.
- [13] Thompson, R. S., E. M. Rantanen, and W. Yurcik, Network Intrusion Detection Cognitive Task Analysis: Textual and Visual Tool Usage and Recommendations, *50th Annual Meeting of the Human Factors and Ergonomics Society (HFES)*, 2006.
- [14] Thompson, R. E-C. S., Usability of Textual and Visual Interfaces for Computer Network Security Engineers, *M.S. Thesis, University of Illinois/Human Factors Division*, Sept. 2006.
- [15] Whitten, A. and J.D. Tygar, Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0, *USENIX*, 1999.
- [16] Yin, X. et al., VisFlowConnect: NetFlow Visualizations of Link Relationships for Security Situational Awareness, *ACM VizSEC/DMSEC*, 2004.
- [17] Yurcik, W., R.S. Thompson, M.B. Twidale, and E.M. Rantanen, If You Can't Beat'em Join'em: Combining Text and Visual Interfaces for Security System Administration, *ACM Interactions*, Jan./Feb. 2007.
- [18] Yurcik, W., Visualization Tools for Security Administrators, *8th Intl. Financial Cryptography Conf., Lecture Notes in Computer Science*, Vol. 3110, Springer-Verlag, 2004.
- [19] Yurcik, W., J. Barlow and J. Rosendale. Maintaining Perspective on Who Is the Enemy in the Security Systems Administration of Computer Networks, *ACM CHI Workshop on System Administrators Are Users, Too*, 2003.
- [20] Yurcik, W. et al., Two Visual Computer Network Security Monitoring Tools Incorporating Operator Interface Requirements, *ACM CHI Workshop on Human-Computer Interaction and Security Systems (HCISEC)*, 2003.