

SESAME: Extending the Desktop Metaphor to Support Security Decision Making

Jennifer Stoll
Georgia Institute of Technology
Atlanta, GA 30332
jstoll@gatech.edu

Craig Tashman
Georgia Institute of Technology
Atlanta, GA 30332
craig@cc.gatech.edu

W. Keith Edwards
Georgia Institute of Technology
Atlanta, GA 30332
keith@cc.gatech.edu

ABSTRACT

We describe an interface design to help non-expert computer users make better informed decisions about their computer security. Our interface design, called Sesame, uses a direct manipulation extension of the ‘desktop metaphor’ and four tightly coupled views to provide users with information pertinent to several types of security threats. We have performed a series of evaluative user studies, and discuss how our current design has been informed by these studies.

1. INTRODUCTION

Given the increasing number and variety of attacks on end-users and the increasing complexity of the security landscape, more security decisions must be made on end-user systems than ever before. However, because no security tool has a comprehensive understanding of the user’s larger context, or the security-convenience tradeoffs she may be willing to make, these decisions still largely fall on the user. Consequently, users must now understand more about their computers and networks to make *intelligent* security decisions. However, when popular consumer security tools alert users to potential threats and require them to make decisions, users are given data that is neither sufficiently detailed nor sufficiently comprehensible to allow them to make those decisions in an informed manner. For example, alerts are commonly of the form, “[Program X] is trying to access the trusted zone”. The information then given to the user to help them with the decision of whether to allow such access is an IP address and a port number. Unfortunately, this is only a fraction of the information pertinent to the threat, and it is not meaningful to most non-expert users.

Given the paucity of the non-expert user’s mental model, the quantity of information of which they must be made aware, and the shortcomings of existing tools, we have developed a novel interface for security tools targeted at non-experts. This interface combines spatial/physical metaphors and information visualization techniques to provide users with a better understanding of complex security information. Rather than just a “prettier” interface, our goal is to provide users with a tool that supports “security learn-ability” [2]. Our interface, called Sesame, provides more comprehensive information necessary for security decision making by opening the metaphoric desktop interface. Currently this interface exists as an early, interactive prototype that runs on the Windows platform; the features described below are largely accessible in this prototype, although it is not currently connected to “live” network or process data.

2. SESAME VISUALIZATION

We have taken a two-pronged approach to the design of Sesame: 1) data representation through tightly coupled views and 2) education through visualization. The first prong is to present several highly integrated views of the data that allow users to identify unusual trends in several different time scales, and subsequently investigate those trends to evaluate whether they represent threats. The second prong is to frame this presentation of the data in such a way that it enables its own interpretation. For this we provide a spatial, direct manipulation extension to the desktop interface where users can observe the significance of unfamiliar abstractions like processes and network connections in terms of familiar abstractions like windows and geographic locations. At a high level, Sesame’s extensions to the desktop metaphor are designed to allow the user to “open” the desktop UI, to observe running processes and network connections and their relationship to the visual elements on the desktop itself.

The information conveyed by Sesame is centered on helping users mitigate several types of threats: spyware, phishing, and bots. Consequently, the visual elements we display are relatively process- and network-centric.

2.1 Representation

While a direct manipulation, metaphor-driven interface is pedagogically valuable, it is not ideal for representing large quantities of data. To get around this limitation, our visualization consists of four tightly coupled, tiled views (Figure 1). Only one of these views represents data through the use of a spatial metaphor. The others represent data in more abstract terms, allowing more information ‘density’, while relying on the first view to teach the user the meaning of the different abstractions.

The direct manipulation interface to which we refer is seen in the first of the four views provided in our tool. When the user initially invokes Sesame, their Windows desktop appears to rotate around the vertical axis revealing cube-shaped objects that are ‘connected’ to the open windows from behind. These cube-shaped objects represent processes or ‘engines,’ as we refer to them. We show the connections from these engines to remote computers, with the latter being represented using the geographic locations of the remote systems. We employ semantic zooming to let users to learn more about engines and remote systems. For example, when the user hovers over an engine with the mouse, the engine will slightly expand to show a better view of the details about the process such as CPU usage and verification information. The user can then click to expand the engine further to view more elaborate

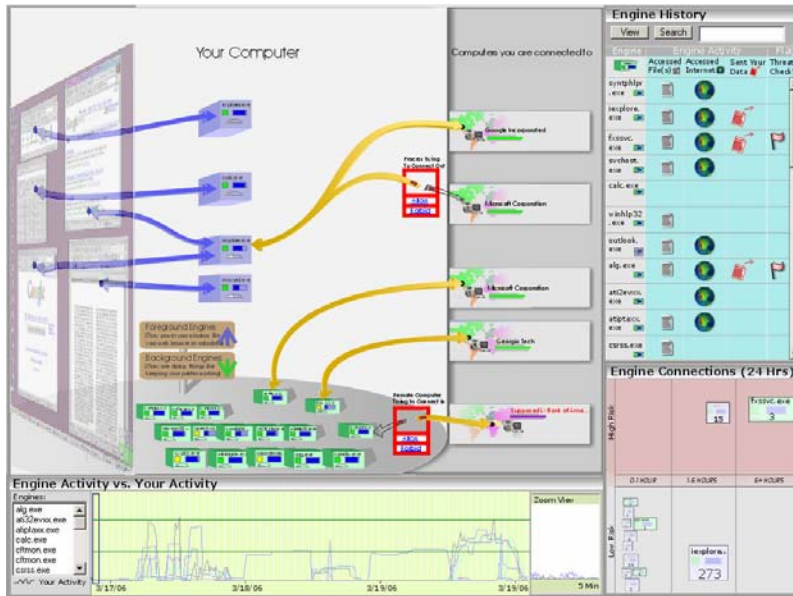


Figure 1. (Clockwise from top-left) Metaphor view, ‘Engine History’ view, ‘Engine Connections’ view, and ‘Engine Activity vs. User activity View.’

details such as policy associated with the process, its age and information about its manufacture.

While the Metaphor View represents the present state, the remaining three views reveal facts about the history of the user’s system. The ‘Engine History’ view provides users with at-a-glance statistics about their engines on a long-term time scale. It lists all processes ever seen in the Metaphor View, and provides details such as whether they have connected to the internet or have accessed the file system. When the user clicks on an item in the list, the item’s row will drop down to show additional facts about the engine, and what, if any, course of action is recommended. The user can also choose the time window for these statistics to be an hour, a day, a week, or a month. Using a behavior-based heuristic, suspicious processes are flagged and the user can click the flagged process to learn the details about its status and past actions.

The ‘Engine Activity vs. Your Activity’ view lets the user compare engine activity against her own usage of the computer, allowing the user to see, for example, what processes are active when she is not using her computer. It displays three days of activity to let the user identify long-term trends, and includes a 5-minute zoom window to find smaller trends as well.

The shortest term historical view is ‘Engine Connections.’ This view helps users identify processes with suspicious network behavior. It shows processes as a scatterplot, with horizontal and vertical positions corresponding to the duration of the process’s longest connection and its evident risk level. The process’s size in the plot corresponds to how much data it has exchanged.

3. ITERATIVE DESIGN

We have employed an iterative design process to help arrive at our current design. We performed user testing on paper prototypes of our earlier designs, one of which was very similar to our current Metaphor View, to determine which representations of system data would be most comprehensible to non-experts. These

tests consisted of task-based trials with our prototypes as well as think-aloud use of them. We asked volunteers how they would perform certain tasks with Sesame, and to explain what they thought Sesame was showing. The results of our user studies motivated large scale modifications in our design. For example, the use of multiple tiled views to show more engine information such as connection history, file accesses, and comparative system activity levels was the result of speaking with study participants as well as security experts. Given the amount of additional data that this, our second design iteration, makes available, we use a layering scheme in presenting data. In keep with the safe-staging [3] principle, we layer the details such that the user can explore and learn more information at her own pace. As one participant aptly expressed, “[Sesame] would teach me then as I’m looking at it, it teaches me the different things that are going on inside the computer. / It would be nice to be able to see [and say] ‘Oh that’s what it is’...”

The most encouraging result of our study however, was that it showed that some users understood how they would use our tool. After very little explanation, such as explaining the purpose of an engine, the

participants described how the concrete representation of connections between windows, processes and the network could help them detect a simple attack such as a phishing scam.

4. CONCLUSION

We have described Sesame, a metaphor-driven security interface design employing techniques of information visualization and direct manipulation to give non-expert users the knowledge they need to make informed security decisions. To do this, we use multiple views representing different time scales. Based on user study results, we believe that non-expert users can understand some of the internal structure of their computer and its relationship to the network through these interface concepts, perhaps leading to better-informed security decisions on the part of non-expert users.

As noted earlier, Sesame currently exists as an early, interactive prototype running on Windows. Our future work includes a further round of user testing to evaluate our current design, followed by refinements to the prototype to support more interactivity and higher fidelity.

5. REFERENCES

- [1] Dourish, Paul, Redmiles, D., An Approach to Usable Security Based on Event Monitoring and Visualization. *News Security Paradigms Workshop '02* (Virginia Beach, VA 2002).
- [2] Sasse, Angela M., Has Johnny Learnt to Encrypt by Now? Examining the Troubled Relationship Between a Security Solution and Its Users. *5th Annual PKI R&D Workshop* (2006). middleware.Internet2.edu/pki06/proceedings/sasse-johnny_usability.ppt
- [3] Whitten, Alma, Tygar, J. Safe Security Staging. *CHI 2003 Workshop on Human-Computer Interaction and Security Systems* (Ft. Lauderdale, FL 2003).