

Security Monitoring, Visualization, and System Survivability: A Position Paper for ISW-2001

Philip E. Varner and John C. Knight
Department of Computer Science
University of Virginia
{philv|jck}@cs.virginia.edu

September 10, 2001

1 Introduction

A significant impediment to the development of large-scale survivable systems is the inability to accurately monitor these systems in real-time. Traditional methods of monitoring rely on system logging and textual messaging to relay information from the system to the administrators. While this works for small systems consisting of at most of a few dozen nodes, its inherently serial human interface does not and cannot scale to national or global proportions. To overcome this impediment, we must create visualization and sonification techniques that will allow us to easily monitor our survivable systems in order to quickly detect and respond to malfunctions or attacks. Here we must remember that survivability is not synonymous with reliability or availability, and that any steps we can take to decrease system degradation are important [4].

The primary step towards adequate security, and therefore survivability, is prevention by properly configuring all devices and implementing a robust network security architecture. However, even the best laid plans can be foiled, whether by advanced or unseen techniques, a known buffer overflow software flaw, or common script kiddie tools. In many systems, attackers are only noticed after they have already breached defenses and possibly caused significant damage. If a system is methodically and consistently monitored, an administrator can recognize attacks as they occur and take action to defend against them.

Some of the problems with prevention outlined in Mukherjee et al. include the impossibility of building an absolutely secure system using current software development techniques and technologies, the impracticality of discarding current open systems (such as the Internet) in favor of new secure systems, the hindrance of overly cautious prevention mechanisms to a user's productivity, the fallibility of cryptographic techniques and methods, and the possibility of abuse by legitimate users [8]. For these and other reasons, we must not only fortify our defenses to prevent security breaches, but we must

actively monitor hosts, network devices, and network traffic to detect, visualize, analyze, and defeat malicious activity.

2 Intrusion Detection

Intrusion detection system (IDS) research was a hot research topic in the late 1980s and early 1990s. Most projects focused on either building rule-based expert systems to detect known system attacks or on statistical anomaly detection to detect deviations from normal system activity. Most of these systems worked well at detecting known attacks in a clean room environment. However, because attacks in the wild are so sophisticated and evolve very rapidly, little useful progress into deployable systems was made. Current research is mainly focused on the application of data mining techniques and mobile agent technologies to the problem.

The challenge of intrusion detection is not increasing sensitivity to intrusions, but decreasing the number of false positives. According to Axelsson, “the false alarm rate is the limiting factor for the performance of an intrusion detection system” [1]. Axelsson argues that effectiveness of an IDS is constrained by the base-rate fallacy. The basic premise behind this is that for large data sets, even very low false alarm rates create a large number of false positives. This indicates that the primary problem we should be addressed in intrusion detection is not simply discovering any possible intrusion, but only discovering intrusions. A possible solution to the problem lies in the application of visualization and sonification techniques to the aggregation of distributed intrusion detection data. Many IDSs rely on one or two network sensors to determine if an attack is occurring. Instead of this, we combine data from *many* sensors and display this data such that a human viewer can derive real meaning from it and quickly comprehend any intrusion activity in the network.

3 Visualization

Humans are inherently visual beings, with over half the brain being dedicated to visual information processing [2]. Visualization can be defined as “a way of making pictures from data that engage processes effectively” [2]. To create a visualization is not merely to form pretty pictures, but to “map the information into a physical space that will represent relationships contained in the information faithfully and efficiently” [3]. The bandwidth of the human visual system is greater than any other sense, allowing humans to see and understand huge amounts of complex data quickly and accurately. A demonstration of this is the ability of a person to glance into a crowd of people and recognize a friendly face. With visual information processing, data is not only processed by the brain faster, but fundamentally changes our processing strategy. Instead of using conscious mechanisms (i.e. I read something, I translate it into a mental model, I understand the mental model), visual processing uses preconscious mechanisms. These mechanisms are “hardwired, highly parallel processes that handle the initial stages of analysis of the retinal patterns” [2]. Visualization moves human data input from an inherently serial process, such as reading text, to the fundamentally parallel process of visual perception, thereby increasing information uptake and understanding.

Computer graphics began in the late 1960s and was primarily concerned with scientific visualization. The use of computers to see existing reality and create new realities was a fundamental shift in information processing. Previously, computers were used simply as a convenient mechanism to store, process, and textually display data, but the advent of computer graphics allowed one to transform this data into a more understandable and intrinsically communicative form of information. The old adage “a picture is worth a thousand words” is quite true, since massive volumes of raw data can be transformed into images that reveal meaningful aspects of the information.

While many people associate visualization only with optical transmission, the primary objective is to transfer data into a mental visualization. So, we include sonification as a possible aid to visualization. Lodha, Joseph, and Renteria find that “bi-modal visual and sound data mappings together provided more accurate understanding of data displays” [6]. Sonification, or mapping of data to non-speech sound, can use parameters of pitch, volume, timbre, duration, frequency, amplitude, and rhythm [5, 7] in a sound clip.

4 Prototype System

To explore the possibilities of such systems, we are developing a prototype that incorporates novel visualization and sonification techniques for IDS data. The general requirements for this system are:

- Flexible - can easily and quickly be adapted to new attacks and new monitoring methods and software. Existing monitoring programs should not require modification to be used with the system.
- Scalable - can be used on an arbitrarily large or small number of hosts. For practical purposes (mainly due to limitations in display technology), we set an arbitrary target at 10,000 hosts.
- Useful - the system will analyze and filter information such that a person or persons looking at the display can easily view it in near real-time and understand events occurring in the network.

The most practical way to fulfill these requirements is to design the system in a modular fashion such that it can be easily extended and configured. Thus, the visualization module is independent of the data source and there is a well defined communication mechanism. This allows us to gather data from many different types of monitors and add new monitors when they are created. In the long run, the system will work generally in the following way:

- Various monitoring programs such as portsentry, snort, Tripwire, EMERALD, scanlogd, and NFR monitor the ports, web server logs, system logs, and network traffic to provide system activity.
- An Agent takes data from these tools and searches for information corresponding to a broadly-defined rule set. The data is securely sent to subscribed entities, known as Aggregators and Watchers.

- An Aggregator receives the Agent's data and inserts it into a database. Watchers may then subscribe to receive data from Aggregators.
- A Watcher can subscribe to multiple Aggregators or Collectors. It takes received data and creates a visualization/sonification of this data. If the amount of data is too great for a single machine to process, a distributed system can be used. Novel display mechanisms may also be employed, including multi-monitor displays, touch-based interaction, display walls, and polyphonic sound arenas.

5 Summary

By combining visualization, sonification, and intrusion detection, we can create large-scale system interfaces that will allow us to more easily monitor these systems in real-time for malicious or dysfunctional behavior. Through useful graphical monitoring, we can more quickly react to undesirable events and therefore increase the survivability of these systems.

References

- [1] Stefan Axelsson. The base-rate fallacy and the difficulty of intrusion detection. *Trans. Inf. Syst. Secur.*, 3(3):186–205, Aug 2000.
- [2] Richard Mark Friedhoff and Mark S. Peercy. *Visual computing*. Scientific American Library, New York, 2000.
- [3] Nahum Gershon, Stuart Card, and Stephen G. Eick. Information visualization tutorial. In *Proceedings of the conference on CHI 98 summary: human factors in computing systems*, pages 109–110, 1998.
- [4] J.C. Knight and K.J. Sullivan. On the definition of survivability. Technical Report CS-TR-33-00, University of Virginia, Department of Computer Science.
- [5] G. Kramer. *Auditory Display, Sonification, Audification, and Auditory Interfaces*. Addison-Wesley, 1994.
- [6] Suresh K. Lodha, Abigail J. Joseph, and Jose C. Renteria. Audio-visual data mapping for GIS-based data: an experimental evaluation. In *Proceedings of the workshop on new paradigms in information visualization and manipulation in conjunction with the eighth ACM international conference on Information and knowledge management*, pages 41–48, 2000.
- [7] Suresh K. Lodha, Catherine M. Wilson, and Robert E. Sheehan. Listen: Sounding uncertainty visualization. In *IEEE Visualization '96*, pages 189–196, 1996.
- [8] B. Mukherjee, L.T. Heberlein, and K.N. Levitt. Network intrusion detection. *IEEE Network*, 8(3):26–41, May-June 1994.