An Immunological Approach to Intrusion Detection

Andrew Watkins
Mississippi State University
Department of Computer Science
PO Box 9637
Mississippi State, MS 39762
(662) 325-2756   fax: (662) 325-8997
andrew@cs.msstate.edu

Abstract

This paper presents an examination of intrusion detection schemes. It discusses traditional views of intrusion detection, and examines the more novel, but perhaps more effective, approach to intrusion detection as modeled on the human immune system. The discussion looks at some of the implications raised by intrusion detection research for information security in general.

Keywords: intrusion detection, computer security, information security, security, biological approaches to computing.

1. Introduction

The increasing complexity of modern computing systems makes traditional views of information security impractical, if not impossible. As pointed out in [5] and [8], computing environments are dynamic with near constant changes in configurations, software, and usage patterns. This makes completely securing a given system—a difficult theoretical task for static systems—unfeasible for the dynamic nature of today's systems. This presents the need for a more dynamic view of information security, one that recognizes the insufficiency of static descriptions of policy and security mechanisms and that proposes a dynamic means of providing security which is sufficient [8] for a given system at a given time.

One of the primary vulnerabilities of today's heavily networked systems is the susceptibility to intrusion. An intrusion attempt can be defined as [7]:

The potential possibility or a deliberate unauthorized attempt to:

- access information,
- manipulate information, or
- render a system unreliable or unusable.

In other words, an intrusion attempt is any threat to the confidentiality, availability, or integrity of the information on a given system.

In general, the focus of this paper is the mechanisms for detecting such intrusions (Intrusion Detection Systems (IDS)). Section 2 gives a brief overview of intrusions and intrusion detection schemes. Section 3 explores the analogy of IDS and information security to the human immune system and argues that the dynamic nature of today's computing environment is better suited to this approach than to traditional views of protection. Section 4 offers a discussion of the implications of IDS research, which is followed by a summary and conclusions in Section 5.

2. Intrusion Detection Schemes

While there are numerous ways of categorizing types of intrusions, IDS are typically divided into two broad categories of misuse detection schemes and anomaly detection schemes [7]. Misuse detection techniques are based on the assumption that, given known patterns of attack, it is possible to detect when these patterns (or, more importantly, attacks) are occurring. Simple examples of this technique include the noticing of repeated failed login attempts or virus scanners, which recognize the "signature" of a virus. While extremely effective for known threats, misuse schemes are virtually useless in detecting novel patterns of attack or unknown threats.

Anomaly detection systems assume that there are patterns of normal behavior for a system. This assumption makes detecting intrusions a matter of comparing the behavior in question to the normal behavior. If there is a significant difference, then it is likely that an intrusion is occurring or has occurred. One example of such a system is the data mining of audit logs for patterns of events that are likely to occur during normal operation. If an event, or sequence of events, occurs that is not predicted by these rules of normal behavior, then it is anomalous and is most likely an intrusion [4].

Problems exist for both misuse and anomaly detection systems, especially when only one technique is used exclusively. Many current IDS tend to be monolithic systems with little or no provision for the failure or subversion of the IDS itself. Obviously, misuse detection systems require prior knowledge of the general type of intrusion likely to occur, and these techniques fail to recognize novel attacks. Anomaly detection systems, while arguably more powerful, are often susceptible to an intruder slowly training the IDS by gradually varying from normal behavior. Eventually, the IDS could be taught to accept inappropriate (threatening) behavior as normal. However, this does not mean that misuse and anomalous detection schemes are useless. Indeed, when used jointly on a more microscopic, diversified level (as opposed to a single monolithic IDS), these two schemes provide a basis for a dynamic, intuitively effective, model for IDS.

3. Computer Immunology

One way of viewing intrusion detection is the task of distinguishing “self” from “non-self.” The self of a system includes all the legitimate functions, processes, or data accesses of the system, and the non-self is everything else. This ability to distinguish self is also found in the human immune system [2], [3], [5].

As mentioned previously, one of the primary weaknesses of traditional approaches to IDS is the use of a single system that monitors only one data source for possible intrusions or, if used in conjunction with various intrusion detection techniques, is often easily subverted. The human immune system does not suffer from this particular weakness. The detectors in the immune system are light-weight, diversified, autonomous agents that work separately, yet are interconnected, to identify (and ultimately remove) those elements in the body that are not recognized as self. The following two subsections explore some of the research that has been undertaken in an attempt to use the model of the immune system as a valid approach to intrusion detection.

3.1 Autonomous Agents

In their report "Active Defense of a Computer System using Autonomous Agents," Crosbie and Spafford propose the use of autonomous agents for intrusion detection [1]. Each of these agents would be highly specialized to monitor only certain activities of the system, such as writing to a particular file, accessing a particular IP address, or even lower level functions. No one agent is critical to the function of the IDS, just as no one agent has the ability to independently declare that an intrusion is occurring. Rather, each agent monitors its particular niche and, upon detecting behavior that is considered suspicious, will send out a signal to the other agents in the system. Having received this signal of suspicion from the first agent, each agent will heighten its sensitivity to (or lower its tolerance of) suspicious behavior. As the suspected intrusion continues into the system, encountering more and more agents, the level of suspicion will eventually exceed a pre-determined threshold, and an agent will alert the system administrator that an intrusion is occurring.

As with all intelligent IDS, the autonomous agent approach requires a significant amount of training while each agent learns what to monitor and what suspicious behavior entails. This requires a knowledgeable security officer to determine not only the proper level of reaction by the agent, but also which resources need to be protected most. Crosbie and Spafford envision that over time the agents will develop techniques for recognizing newer threats to the system, but before the agents are released, they must be trained concerning the policies and potential vulnerabilities for the given computing environment. Thus, the use of autonomous agents involves both misuse and anomaly detection techniques.

3.2 Distinguishing Self

While there are a wide range of characteristics the human immune system could have examined to detect an intruder in the body, the immune system uses numerous detector cells (T-cells) which, based on the protein sequences of the intruder, bind with that intruder and eliminate it[*]. One of the difficulties of this method is that the intruder and the host body are both composed of similar proteins, so the detector cells must be able to distinguish self from non-self. This is similar to the difficulty of detecting an intrusion in a computer system. The processes of the intruder are similar in nature to the legitimate host processes. The T-cell concept can be applied to the computing environment by the generation of detectors (somewhat similar to autonomous agents) which monitor the system calls made by various processes and are able to distinguish a legitimate process (self) from an illegal one (non-self). Researchers at the University of New

[*] For a more detailed discussion of the immune system in general and how this relates to computer systems see[2, 5].

Mexico [3] found that processes (such as sendmail, lpr, etc.) produce distinct sequences of calls to the system. The detectors monitor these calls, and if they vary from the known normal calls of the process, then an intrusion is occurring. Additionally, it was found that a window of only 6 system calls was sufficient to distinguish self from non-self. This concept is not just limited to operating system calls. Stillerman, Morceau, and Stillman have had some success in using detectors in distributed applications as well [6].

Some of the advantages of this view of intrusion detection are that it is highly distributable and that individual versions of software or system configurations can generate unique sets of detectors. So, while an intruder could possibly subvert an individual detector, the intruder would not necessarily gain access to the entire host and would most certainly be prevented from unwarranted access to an entire network. Due to the computational costs of protecting every process in the system, this view of IDS will necessarily be less comprehensive than the more traditional schemes, but it also prevents an intruder from running rampant throughout a system. This concept of intrusion detection, taken a step further, would allow the detectors not only to notice an intrusion, but also to stop the intrusion either by severing the given connection or by isolating (and perhaps terminating) the compromised program or process until the system administrator can be notified of the intrusion.

4. Discussion

The analogy of intrusion detection as an immune system allows for a fairly radical shift in the way security is viewed. The immune system is necessarily dynamic, as it must be resistant to both the old and new intruders it faces. While each person’s immune system is similar, no one person's set of detectors is exactly the same as another's. This enables a given intruder to be thwarted more easily. The direct implication to a networked computing environment is easy to see. If each host in a given network contains a different set of detectors protecting it, then intrusion into one host will most likely be limited to that host. Or, on a lower level, penetration of one process on a system will be stopped at that process. Since the detection mechanisms are small and numerous, the subversion of one agent does not necessarily imply the compromise of the entire IDS. This could allow for cheap renewal of the IDS, and this makes protecting against novel threats more easily accomplished.

Ostensibly this has been a discussion about intrusion detection systems. However, IDS research raises some interesting implications for information security policy. In some ways, the challenge of developing a viable intrusion detection system is the same challenge of information security in general: How can an IDS be designed to accurately distinguish legitimate behavior from an intrusion? How can an IDS be designed to recognize known and unknown threats? Like all mechanisms in a computing environment, policy plays a large role for a given IDS. The security policies for a given system will directly determine what is considered anomalous and what behavior is within normal bounds. Policy also points to areas of known weakness and expects those areas to be protected.

Vaughn proposes a cyclical process to providing sufficient information security for a given computing environment at a given time [8]. This process is considered continuous and must provide security that is flexible enough to change as threats to the system change. IDS can play a key role in this flexibility. In general, IDS (particularly those with some amount of anomaly detection) can help to identify weaknesses in the system. Awareness of these weaknesses can then be used during the process of updating the information security policies and procedures for the environment in question. So, there can be a certain amount of mutual dependency between an IDS and the policies of the system to be protected. Naturally, IDS is only one mechanism that can provide appropriate feedback for the periodic updates of security policies. However, IDS can potentially provide a great deal of guidance in the fine-tuning of policy.

5. Summary and Conclusions

Both misuse and anomaly detection schemes are valuable for detecting specific intrusions into a given system, but are not necessarily dynamically suited for today's detection needs. Misuse detection techniques are ideal for recognizing and preventing known threats; however, this is not helpful when faced with an unknown or new threat to the system. Anomaly detection systems, by creating a signature for normal behavior and recognizing abnormal behavior as an intrusion, are more dynamic, but these systems can be trained to recognize abnormal behavior as normal with time. Therefore, an IDS that can function using both misuse detection and anomaly detection would be ideal.

Combining elements of both detection techniques would create an IDS that can not only recognize abnormal behavior, but that can learn from this behavior to prevent intrusion. This system would begin with the anomaly detection system by targeting deviations from normal behavior and recording these deviations. Then the misuse detection system would be able to recognize the signatures of the deviations as they occur, and, after a set number of deviations have been made, would signal that an intrusion was occurring.

Similar to the human immune system, this dynamic IDS would be able to recognize threats to any part of the system and learn from these threats to make future recognition more expedient. By partitioning the different processes of the operating system or network with each process or application having its own intrusion detection agent or detector "cell," the entire system would be alerted when a deviation from normal behavior occurred in any area of the system. Therefore, the IDS would heighten its sensitivity throughout to protect the computing environment as a whole. In other words, by protecting the individual parts of the system, the whole system is, in turn, protected. As intrusions occur, the system would learn to recognize them, creating a more efficient intrusion detection and (hopefully) prevention system. In all, this IDS would be capable of dynamically detecting intrusion while creating reliable predictions of threatening behavior.

6. Acknowledgments

This work is the result of the encouragement of two individuals, Drs. Susan Bridges and Ray Vaughn, whom I wish to acknowledge. Dr. Bridges has provided invaluable support throughout all of my course work and research endeavors. Dr. Vaughn introduced me to the world of information security and greatly aided me in the submission of this paper.

7. References

[1] Crosbie, Mark and Gene Spafford. 1995. *Active Defense of a Computer System using Autonomous Agents.* Purdue University Department of Computer Science. Technical Report 95-008.

[2] Forrest, Stephanie, Steven A. Hofmeyr, and Anil Somayaji. 1997. Computer Immunology. *Communications of the ACM* 40(10) (October, 1997): 88-96.

[3] Forrest, Stephanie, Steven A. Hofmeyr, Anil Somayaji, and Thomas A. Longstaff. 1996. A Sense of Self for Unix Processes. In *Proceedings 1996 IEEE Symposium on Security and Privacy* in *Oakland, California, May, 1996* by IEEE Computer Society, 120-128. Los Alamitos, CA: IEEE Computer Society.

[4] Luo, Jianxiong. 1999. Integrating Fuzzy Logic with Data Mining Methods for Intrusion Detection. M.S. Thesis, Mississippi State University.

[5] Somayaji, Anil, Steven Hofmeyr, and Stephanie Forrest. 1997. Principles of a Computer Immune System. From *New Security Paradigms Workshop* (Presented September, 1997). Downloaded from http://www.cs.unm.edu/~immsec/papers.htm on 15 July 1999.

[6] Stillerman, Matthew, Carla Morceau, and Maureen Stillman. 1999. Intrusion Detection for Distributed Applications. *Communications of the ACM* 42(7) (July, 1999): 62-69.

[7] Sunderam, Aurobindo. 1996. An Introduction to Intrusion Detection. Downloaded from http://www.cs.purdure.edu/coast/archive/data/categ24.html on 15 June 1999.

[8] Vaughn, Rayford B. 1996. A Practical Approach to Sufficient Infosec. *National Information System Security Conference* in *Arlington, Virginia, October 1998*, sponsored by *National Institute of Standards and Technology* and *National Computer Security Center*. Volume 1: 1-11.