

Dagstuhl Extended Abstract: Optimistic Access Control and Anticipatory Forensic Analysis of Insider Threats

Sean Peisert, Matt Bishop, and Christian W. Probst

August 23–26, 2010

Position

Insiders are inside our network, by definition. For the most part, we cannot stop them from taking harmful actions that they are permitted to do by virtue of the access granted by their position without preventing legitimate users from doing their jobs. Our best defense is usually analyzing what they have been doing in the past. Current techniques for post mortem analysis are limited by several hard problems: how does one know what to look for? How does one know what questions to ask of an analysis tool and how to frame those questions? Are current analysis tools capable of answering such questions? We propose a possible set of solutions.

Background

Forensics can be done in an ad hoc approach or through the use of a model, such as *Laocoön* that uses pre-conditions and post-conditions of attack graphs to fit events to attacker goals (targets) or defender goals (security policies) [PBKM07a, Pei07]. The disadvantage of this approach is identifying the goals: security policies can be reverse engineered to a certain extent [BP06, Pei07 §8.2, PH08, PH09], but determining an exhaustive list of possible attacker goals, particularly for a general-purpose computer system, is infeasible, particularly when attacker goals may not be something relating to traditional security concepts of availability, confidentiality, and integrity, but may just be some non-technical goal like embarrassment. Thus, such attack graphs could end up being just as ad hoc or incomplete as existing methods.

Our Premise

We propose a different solution to avoid having to pre-define attacker goals using “guesswork” and without either exposing the system to too much risk or preventing authorized users from doing their jobs. We do this using a multi-pronged approach: first, we present a mechanism that can be used as a starting point for analyzing and comparing actions that a user is predicted to take as opposed to actions that a user actually takes. Second, it can *optimistically* (but cautiously) constrain a user’s future actions when we see that the most likely paths based on execution history are not being followed [BEF10].

Machine learning techniques and hidden Markov models have been applied to intrusion detection for over a decade [WFP99]. However, these previous approaches addressed a somewhat different problem than that which we are interested in: they addressed a binary decision of whether something was anomalous or not—one or zero. The problems that we seek to address are a post-mortem look into the past to see where a particular set of steps deviated from normal; second, a prediction into the future of an attacker’s next likely steps; and third a method of limiting damage. Similarly to past approaches that have involved analyzing function calls for forensics [PBKM07b]. In our case, rather than making a binary decision on whether the sequence appears in the known corpus (or a similar sequence, based on a Hamming distance measurement), our approach uses a probabilistic measurement using *ε-machines*—a special-purpose type of hidden Markov model to analyze the most likely paths based on past usage [BWC10]. Traversal of a less-commonly used path suggests potential real-time observation or post-mortem analysis.

The mechanism works as follows: a corpus of past actions for each user is built over time by monitoring their system call usage. That corpus includes not just statistical information about which

system calls occur but also the order in which they occur. Over time, this corpus defines a weighted graph showing the paths that a user has historically taken, with the largest weights assigned to the most common paths. When a user takes a path that appears more rare (that is, below a certain arbitrary threshold of frequency), optimistic countermeasures are imposed, such as limiting access (e.g., granting read-only, but not write, access) until the countermeasures are countermanded, for example by a security officer and trust is therefore increased because the actions are determined not to be those of an attack with heightened access. Alternatively, a user can be asked to re-authenticate, and restrictions could be optimistically and temporarily overridden unless additional security thresholds are triggered. Karger described an intelligent naming subsystem [Kar87], which reports suspicious events to the user and asks what to do. In our system, this can be extended by applying the scheme to optimistic access control. The system contains a mechanism where supervisor is notified when someone goes down a path of low probability; the supervisor can allow or disallow it (in either case, full logging continues). Note the potential for infinite regression, because supervisors can in turn be checked by their own supervisors, etc. How to manage this is an open research question.

Finally, if a security violation is suspected, a forensic analyst can then be involved to examine the suspicious paths and determine how likely the deviation was to have been caused by changing job functions or other non-malicious causes such as a detection threshold is set too high. Now forensics becomes a matter of determining why the deviant path was taken, which moves it from the realm of technology back to the realm of people. This has the advantage of building on centuries of experience and interpretation, taking into account context external to the system—something an examination of the system will not provide.

References

- [BEF10] Matt Bishop, Sophie Engle, Deborah A. Frincke, Carrie Gates, Frank L. Greitzer, Sean Peisert, and Sean Whalen. A Risk Management Approach to the ‘Insider Threat’. In Christian W. Probst, Jeffrey Hunker, and Matt Bishop, editors, *Insider Threats in Cybersecurity*, Advances in Information Security Series. Springer Verlag, Berlin, September 2010.
- [BP06] Matt Bishop and Sean Peisert. Your Security Policy is *What??* Technical Report CSE-2006-20, University of California at Davis, 2006.
- [BWC10] Matt Bishop, Sean Whalen, and James Crutchfield. Hidden Markov Models for Automated Protocol Learning. In *Proceedings of the 6th International ICST Conference on Security and Privacy in Communications Networks (SecureComm)*, Singapore, September 7–9, 2010.
- [Kar87] Paul A. Karger. Limiting the Damage Potential of Discretionary Trojan Horses. In *Proceedings of the IEEE Symposium on Security and Privacy*, Oakland, CA, 1987.
- [PBKM07a] Sean Peisert, Matt Bishop, Sidney Karin, and Keith Marzullo. Toward Models for Forensic Analysis. In *Proceedings of the Second International Workshop on Systematic Approaches to Digital Forensic Engineering (SADFE)*, pages 3–15, Seattle, WA, April 2007.
- [PBKM07b] Sean Peisert, Matt Bishop, Sidney Karin, and Keith Marzullo. Analysis of Computer Intrusions Using Sequences of Function Calls. *IEEE Transactions on Dependable and Secure Computing (TDSC)*, 4(2):137–150, April–June 2007.
- [Pei07] Sean Philip Peisert. *A Model of Forensic Analysis Using Goal-Oriented Logging*. PhD thesis, Department of Computer Science and Engineering, University of California, San Diego, March 2007.
- [PH08] Christian W. Probst and René R. Hansen. An extensible analysable system model. *Information Security Technical Report*, 13(4):235–246, Elsevier, 2008.
- [PH09] Christian W. Probst and René R. Hansen. Analysing Access Control Specifications. In *Proceedings of the Fourth International IEEE Workshop on Systematic Approaches to Digital Forensic Engineering*,

pages 22–33, 2009.

- [WFP99] Christina Warrender, Stephanie Forrest, and Barak Pearlmutter. Detecting Intrusions Using System Calls: Alternative Data Models. In *Proceedings of the 1999 IEEE Symposium on Security and Privacy*, pages 133–145, Oakland, CA, 1999.