Prioritizing Intrusion Analysis Using Dempster-Shafer Theory

Loai Zomlot
Kansas State University
lzomlot@ksu.edu

Sathya Chandran Sundaramurthy
Kansas State University
sathya@ksu.edu

Kui Luo
Kansas State University
kuiluo@ksu.edu

Xinming Ou
Kansas State University
xou@ksu.edu

S. Raj Rajagopalan
HP Labs
raj.rajagopalan@hp.com

Intrusion analysis and incident management, *i.e.* the process of combing through IDS alerts and audit logs to identify and remediate true successful and attempted attacks, remains a difficult problem in practical network security defense. The major root cause of this problem is the large rate of false positives in the sensors used by IDS systems to detect malicious activities. IDS systems are currently unable to differentiate nearly certain attacks from those that are merely possible, reducing the value of the alerts to an administrator. Standard Bayesian theory has not been effective in this regard because of the lack of good prior knowledge. This paper presents an approach to handling such uncertainty without the need for prior information, through the Dempster-Shafer (DS) theory that uses a generalization of probabilities called beliefs to characterize confidence in evidence in support of a given hypothesis. DS theory also provides a calculus to compute the level of belief in an aggregate of evidence. We address a number of practical but fundamental issues in applying DS to intrusion analysis, including how to model sensors' trustworthiness, where to obtain such parameters, and how to address the lack of independence among alerts. We present an efficient algorithm for carrying out DS belief calculations on an IDS alert correlation graph, so that one can compute a belief score for a given hypothesis, *e.g.* a specific machine is compromised. The belief strength can be used to sort incident-related hypotheses and prioritize further analysis by a human analyst of the hypotheses and the associated evidence. We have implemented our approach for the open-source IDS system Snort and evaluated its effectiveness on a number of data sets as well as a production network. The resulting belief scores were verified through both anecdotal experience on the production system as well as by comparing the belief rankings of hypotheses with the ground truths provided by the data sets we used in evaluation, showing thereby that belief scores can be effective in taming the high false positive rate problem in intrusion analysis.