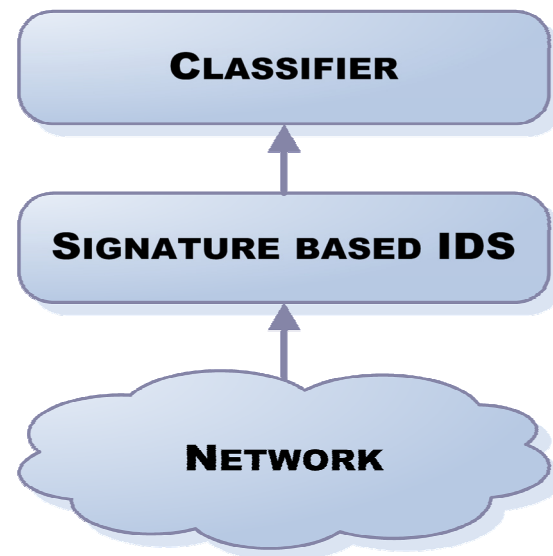




Introduction

The Internet is rapidly becoming a more hostile environment. Different viruses, worms and an increasingly number of hacker attempts are challenges network administrators have to deal with on a daily basis. An important tool to detect misbehaving hosts and network connections, are rule based intrusion detection systems (IDS). One flaw in these types of applications in general, is the number of false alerts they are likely to generate.



Large number of false alerts makes the analyst job, which is an inevitable part of using these systems, a fair bit of a challenge. Chances are that an analyst very well could overlook a true alert, during a spell of large number of false alerts.

A solution which could prevent, or at least reduce false alerts in rule based intrusion detection systems, would therefore be very much appreciated.

Multi-tier intrusion detection system

We introduce a solution to reduce intrusion detection alerts which we call a “multi-tier intrusion detection system”. By this we mean to combine a rule based IDS with a machine learning classification method. Our idea is that a machine learning method could be used to distinguish false intrusion alerts from true alerts with a high rate of correctness. The concept of machine learning, which is currently gaining acceptance in the Internet security world, is to use knowledge of the past to foresee similar events in the future. Often one would call this knowledge training examples. In our case, the training examples consist of previously classified IDS alerts.


The main objective in designing such a system is to reduce the amount of false alerts. However, reducing false alerts would be without any use, if this would mean overlooking a considerable amount of true IDS alerts. Therefore our goals are reducing false alerts and in the same time, not overlook more than a few true alerts.

In our “multi-tier intrusion detection system”, we decided to use the machine learning algorithm naïve Bayesian classifier, an algorithm proven successful in filtering emails and diagnostics. This method has a strength by its general performance in handling a wide set of various training examples. This is typically the case in

classifying IDS alerts. Our solution succeeds according to both these criteria’s. Using training examples in form of logged IDS alerts, we successfully manage to classify a significant number of alerts.

Results and performance

The first tests of our “multi-tier intrusion detection system” showed that we could classify almost 90% of the alerts correctly, but that as many as 14% of the true alerts would be overlooked. Even though the latter result was not as good as we would appreciate, the 90% correct classification is enough to conclude that our method is up for the task considering numerous of sources of noise in the our training examples.



An improved version of our system that included a security margin which main task is to reduce the number of overlooked true alerts, improved the preliminary results considerably. The number of overlooked real alerts was reduced from a total of 14% to the marginal 1.4%. This is a major boost in performance taking security into consideration. Although the effect in detecting

false alerts is 50%, this system could reduce the amount of false alerts to one half of the total, practically without overlooking true alerts.

Main Conclusions

With our results we have shown that a “multi-tier intrusion detection system” could be a tool for reducing false alerts in rule based intrusion detection systems. Combining strengths from to different technologies, “machine learning” and “intrusion detection”, we have implemented a system which can classify IDS alerts and still maintain an acceptable level of security, and by reducing the total of overlooked true alerts to a total of 1.4%, still be able to detect one out of two false alerts. This would be a considerable help in handling such events, letting an analyst concentrate on the more important true alerts.

Further reducing the chance of overlooking particular major security threats, our system could be tuned to work only on alerts defined as minor risk alerts. In this case, a fully automatic solution which completely leaves the alert handling of minor threats to our classifier, could further improve the general performance in the analyse process, leaving considerable fewer alerts for manual analysis. Another approach would be to use a system as we have described as a pure guidance tool. The original version of our tool would be right almost 90%, and could be developed into an advisor system with even better accuracy. This could also turn out to give an analyst a better chance to focus on the high risk attacks.