Foundations of Computer Security

Lecture 73: Intrusion Detection

Dr. Bill Young
Department of Computer Sciences
University of Texas at Austin

Intrusion Detection

An *intrusion detection system* (IDS) can analyze traffic patterns and react to anomalous patterns. However, often there is nothing apparently wrong but the volume of requests.

Note that an IDS is inherently reactive; the attack *has already begun* when the IDS acts.

Intrusion Detection Errors

There are two types of errors when considering any intrusion detection system.

False negatives: a genuine attack is not detected.

False positives: harmless behavior is mis-classified as an attack.

Which do think is a bigger problem?

An intrusion detection system is:

accurate: if it detects all genuine attacks;

precise: if it never reports legitimate behavior as an attack.

It is easy to make an IDS that is either accurate or precise! *Why?* It's hard to do both simultaneously.

Intrusion Detection Errors

An undetected attack might lead to severe problems. But frequent false alarms can lead to the system being disabled or ignored. A perfect IDS would be *both accurate and precise.*

- Statistically, attacks are fairly rare events.
- Most intrusion detection systems suffer from the *base-rate fallacy*.

Base-Rate Fallacy

Suppose that only 1% of traffic are actually attacks and the detection accuracy of your IDS is 90%. *What does that mean?*

- the IDS classifies an attack as an attack with probability 90%
- the IDS classifies a valid connection as attack with probability 10%

What is the probability that a connection flagged as an attack is not really an attack, i.e., a false positive?

There is approximately 92% chance that a raised alarm is false.

Lessons

- False negatives and false positives are both bad for an IDS.
- An IDS must be very accurate or suffer from the base rate fallacy.
- An IDS with too many errors becomes useless.

Next lecture: Anatomy of an Attack: CodeRed