Anomaly Detection

Steven M. Bellovin
https://www.cs.columbia.edu/~smb

Matsuzaki ‘maz’ Yoshinobu
<maz@iij.ad.jp>

Why Anomaly Detection?

- Signatures defend against *known* attacks
  - You need a separate signature for each one
  - By definition, there are no signatures for things that don't exist
- Anomaly detectors look for unusual activity: things that normally don't happen
- Implication: must first know what is normal
  - "Normal" is different for every organization

What's An Anomaly?

Normal

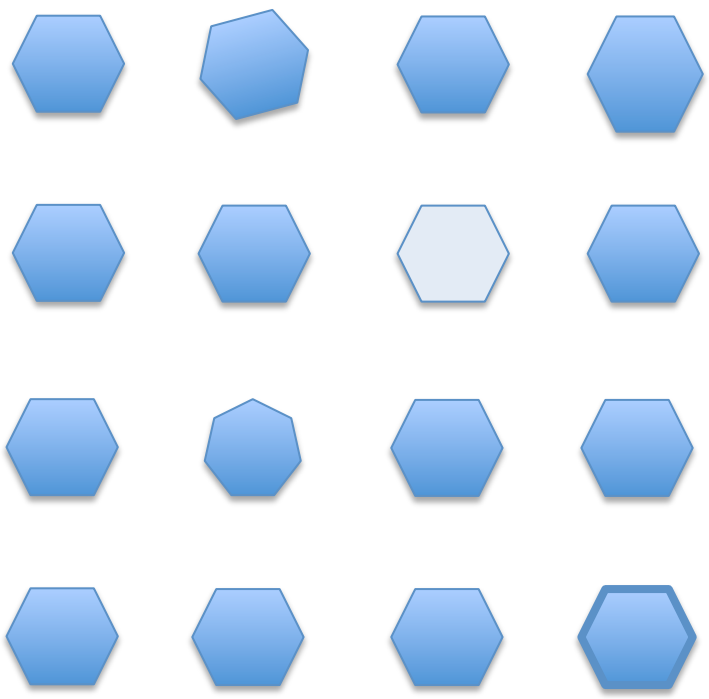

Infected

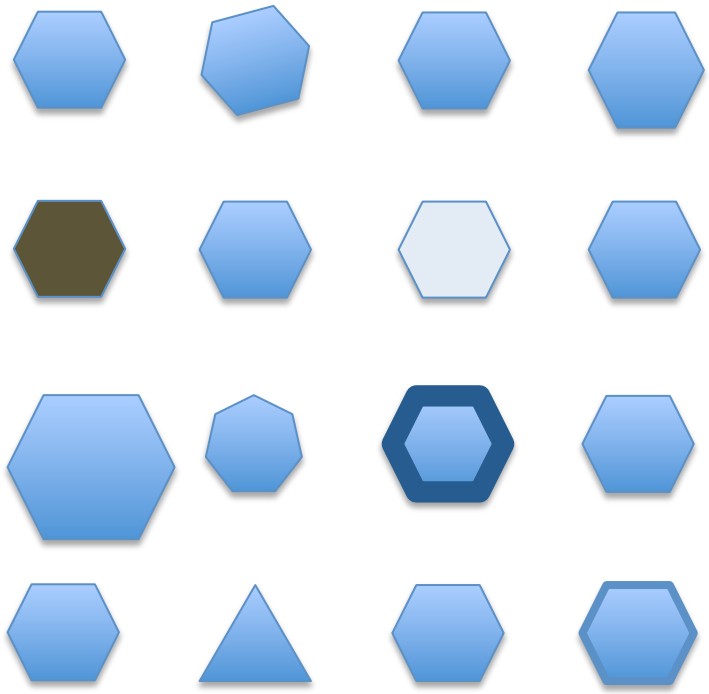

Examples

- massive incoming traffic

  - periodic security update ☺ or DoS ☹

- unusual outbound traffic

  - video chat ☺, flood attack ☹ or information theft ☹

- unusual protocol communication

  - new application ☺ or compromised host ☹

General Process

- Establish a baseline of normal activity
  - Sample activity from times when you're not under attack
- *Train* your detectors on this baseline set
- Continually match current behavior against the baseline
- Investigate "significant" deviations

CV5: “Correlate Violations of Volume, Velocity, Values, Vertices”

- Correlate is obvious
- Violations implies some "normal" model is violated
- Volume and Velocity are standard metrics of expected flow behavior (think highways)
- Values pertain to any content analysis, packet heads, datagrams, email bodies, URL, PHP variable argument values, etc.
- Vertices pertains to graphic theoretic constructs, connectivity between entities, IP addresses, MAC addresses, ports, etc.

Establishing a Baseline

- Different strategies for different uses and kinds of attacks
  - What does your traffic flow normally look like?
  - What applications do users run?
  - What is the byte value distribution of certain file types?
    - Word documents infected with shell code will have more bytes that look like x86 machine code
- Different groups will have different normal behavior

One Way to Define Normal

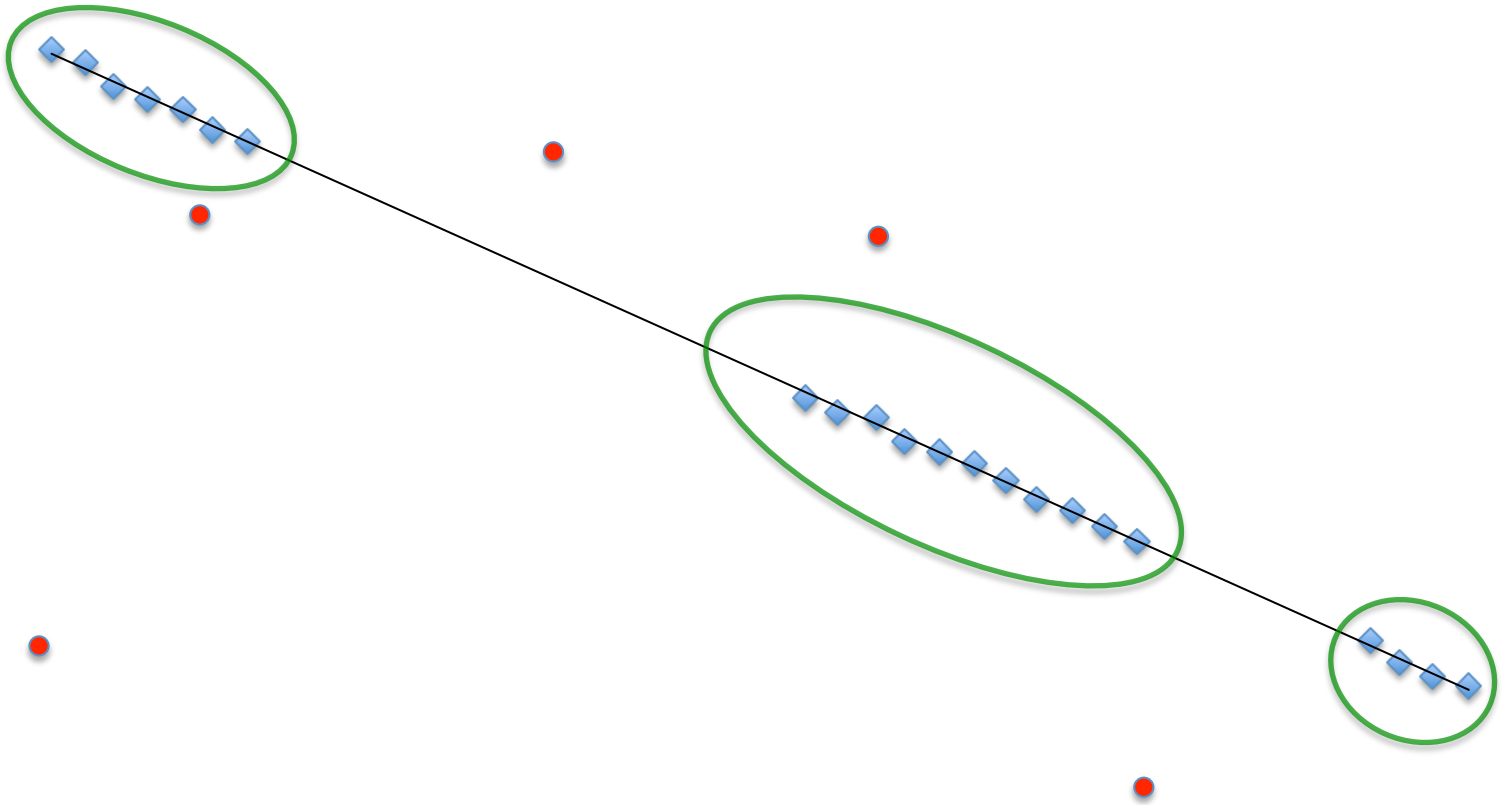

(Mathematically) find clusters. Points outside the clusters are abnormal.

Limitations

- It's hard to define "normal"
  - Was your training data really attack-free?
  - What if legitimate patterns change? New employees? New versions of applications?
- Relatively high false positive rate
- Can miss subtle attacks
- Must run anomaly detectors on many different activities

Advantages

- Can detect minor variants of existing attacks (a serious issue in the anti-virus world)
- Can detect 0-day attacks
- No need to constantly update signature database
- Probably the wave of the future in intrusion detection

Example: Netflow

- router can export traffic flow information (incoming interface, packet headers) to a collector
  - useful to analyze traffic

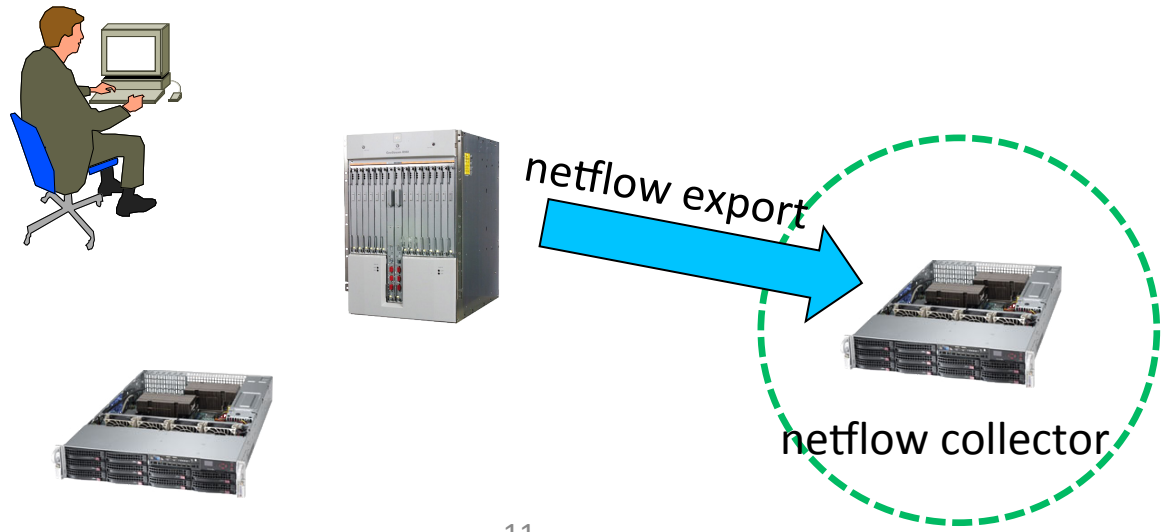

Example: Mail Logs

- Look at the mail logs every day
  - Is someone sending significantly more mail than they normally do?
  - Is someone sending to many more recipients than normal?
  - Is the size of someone's mail messages larger than normal?
  - Anomalies can be benign: recently, someone emailed me a 9 MB, 1600 page PDF, with many scanned images—and it was perfectly legitimate

Example: Host Monitoring

- Monitor system calls
  - What system calls does an application normally make?
  - What sequences of system calls does it normally make?
- Works before encryption or after decryption
- But—attackers can look for and disable a host-based IDS