

Dealing with False Positives in Intrusion Detection


Klaus Julisch, ZRL

kju@zurich.ibm.com

Presentation Overview

- Introduction & Problem Statement
- Characteristics of IDS alarms
- Filtering False Positives
- Problems with Filtering
- Benefits of Filtering -- Looking at What's Left Over
- Conclusion & Open Problems

Introduction & Problem Statement



Presentation Overview

- Introduction & Problem Statement
- Characteristics of IDS alarms
- Filtering False Positives
- Problems with Filtering
- Benefits of Filtering -- Looking at What's Left Over
- Conclusion & Open Problems

Characteristics of IDS Alarms

- Key observation: 5 / 90 Rule

big.gov	small.gov	bank	transport	insurance
80 (81.1%)	153 (83.9%)	47 (43.4%)	15 (50.6%)	298 (41.1%)
66 (7.3%)	47 (7.5%)	56 (15.0%)	80 (20.1%)	55 (18.6%)
9 (6.7%)	80 (3.8%)	80 (12.8%)	47 (12.7%)	119 (13.5%)
47 (1.2%)	15 (3.2%)	15 (8.1%)	29 (6.2%)	83 (8.8%)
29 (1.2%)	35 (0.5%)	66 (7.1%)	321 (1.8%)	349 (5.4%)

Alarm frequencies for five different sensors in April

Presentation Overview

- Introduction & Problem Statement
- Characteristics of IDS alarms
- Filtering False Positives
- Problems with Filtering
- Benefits of Filtering -- Looking at What's Left Over
- Conclusion & Open Problems

Filtering False Positives (1/4)

Filtering by Means of NetRanger's Context Field

Recld	Time	SrcIp	SrcPrt	...	Context
-------	------	-------	--------	-----	---------

NetRanger's Alarm Format

- Idea: Use the context field to verify NetRanger's analysis

Filtering False Positives (2/4)

Filtering by Means of NetRanger's Context Field:
Alarm 80

IIS WEB Server



Black Hat

Get /script.asp%2E?



Problem

big.gov has 126860
type-80 alarms in
April!

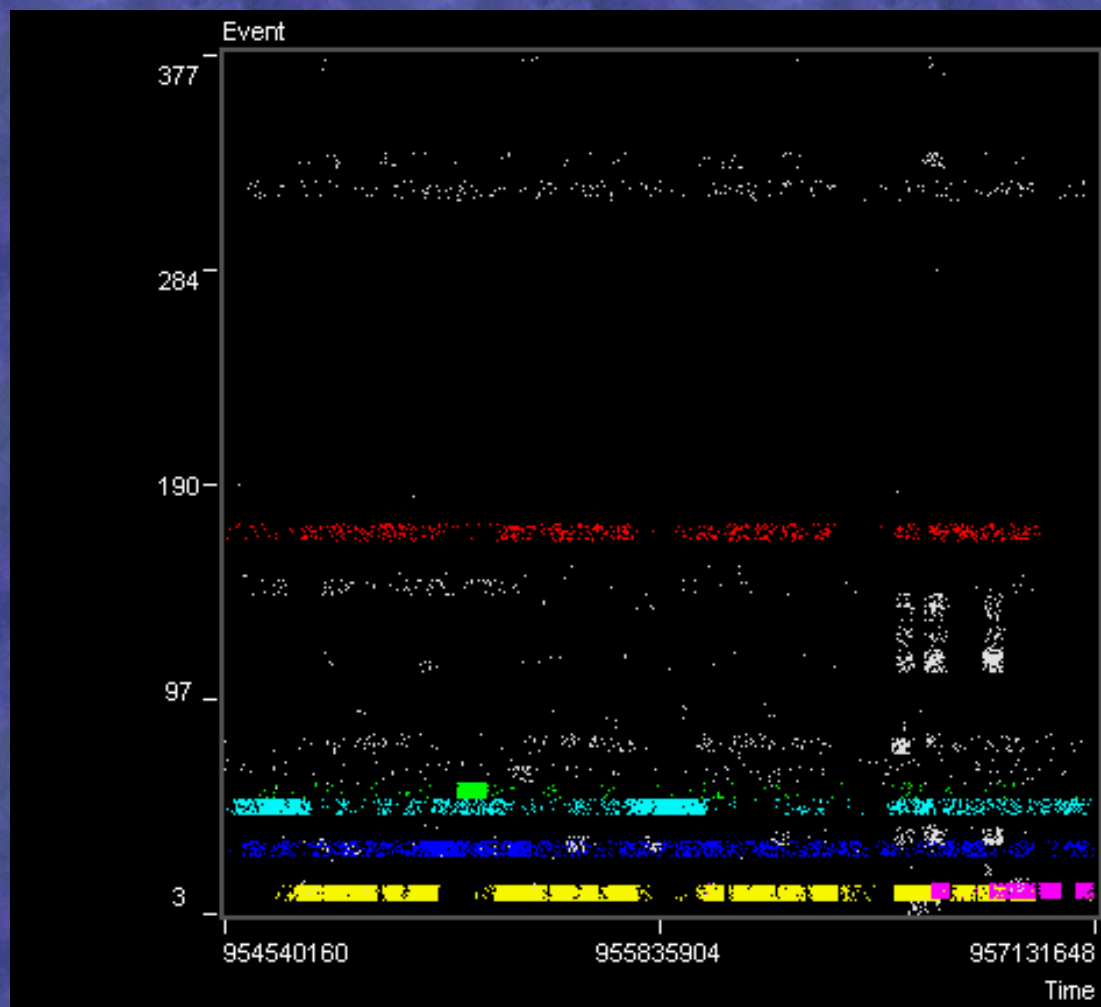
Filtering False Positives (3/4)

Filtering by Means of Alarm Patterns

- Key idea:
 - ▶ Find alarm patterns,
 - ▶ understand their root cause, and
 - ▶ if non-malicious, use the alarm patterns for filtering
- Examples
 - ▶ Host triggers alarm 120 every 30 minutes
 - ▶ Host triggers alarm 65 every $n \times 11$ minutes
 - ▶ Fragmented IP originating weekdays from WEB

Filtering False Positives (4/4)

Fragmented IP Originating Weekdays from WEB Server



Presentation Overview

- Introduction & Problem Statement
- Characteristics of IDS alarms
- Filtering False Positives
- Problems with Filtering
- Benefits of Filtering -- Looking at What's Left Over
- Conclusion & Open Problems

Problems with Filtering (1/4)

Overview

- Finding filtering rules
- Time constancy of filtering rules
- Other problems:
 - ▶ Genericity of filtering rules
 - ▶ Risk of filtering out true positives


Problems with Filtering (2/4)

Finding Filtering Rules

- Not all IDSs provide a "context field"
- NetRanger doesn't provide the context for all alarms
- Finally, how do you find usable patterns
 - ▶ either in the context field, or
 - ▶ in the alarm stream?

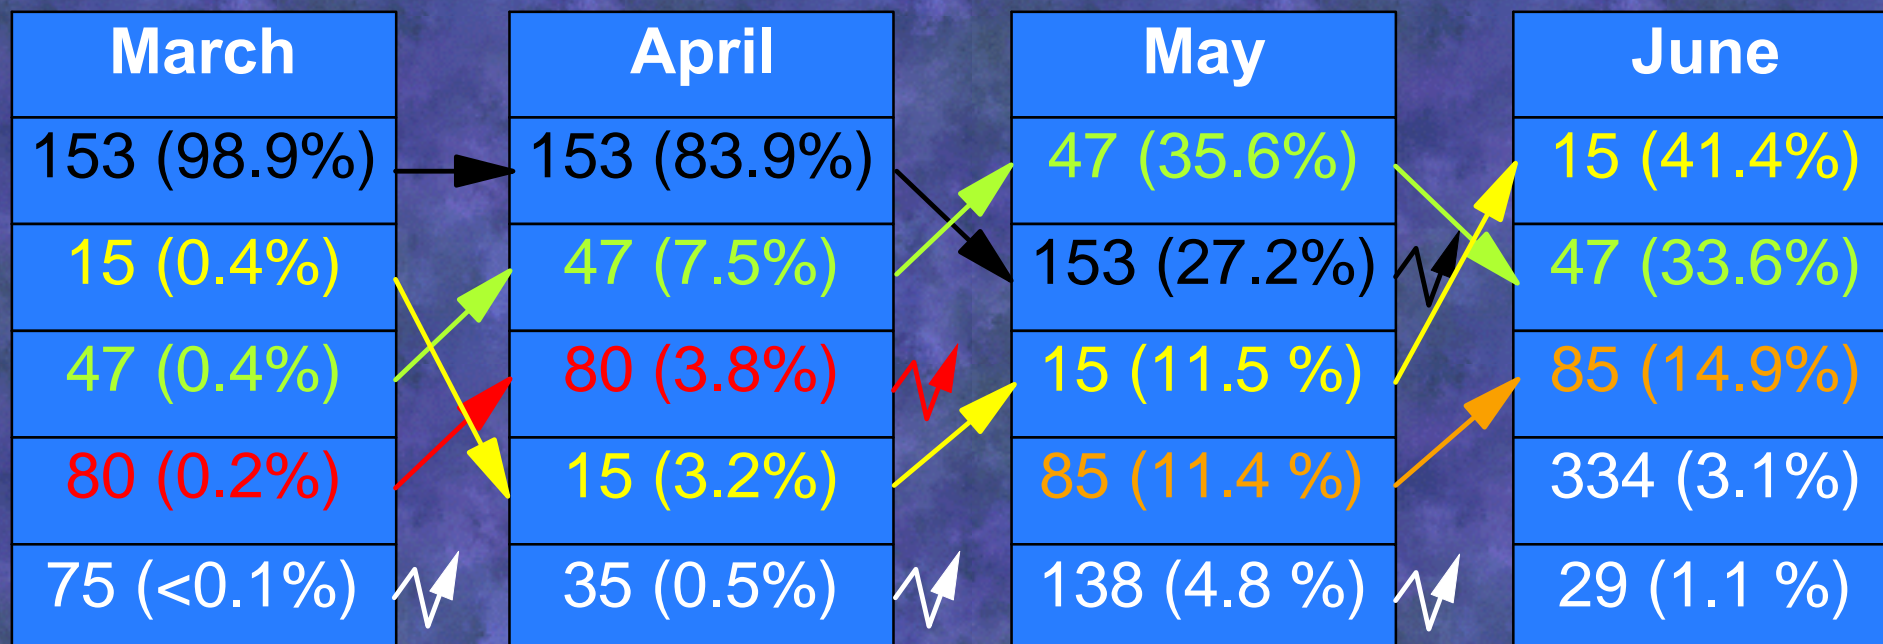
Problems with Filtering (3/4)

Time Constancy of Filtering Rules

- Two aspects of time constancy
 - ▶ Effectiveness of filtering rules over time ✓
 - Alarm 80: 100% - 100% - 100% - 100%
 - Alarm 9: 90% - 79% - 72% - 73%
 - ▶ Justifiability 
 - The most frequent alarms of one month might become insignificant in following months!

Problems with Filtering (4/4)

Time Constancy of Filtering Rules: Justifiability



Development of alarm frequency for small.gov

Presentation Overview

- Introduction & Problem Statement
- Characteristics of IDS alarms
- Filtering False Positives
- Problems with Filtering
- Benefits of Filtering -- Looking at What's Left Over
- Conclusion & Open Problems

Benefits of Filtering (1/3)

The Trace of an Attack Tool



36

burst of {20, 34, 110, 111, 115}

burst of {121, 123, 124}

burst of 133


burst of 138

...

Pause


Benefits of Filtering (2/3)

Coordinated Attack: Multiple Sources Attacking one Target



Start time: April 13, 00:35

synchronized alarms



Target



Benefits of Filtering (3/3)

Opportunistic Attack: Wide Scan



Time Period: 3 days



evil.com

One
probe/target



A.net1



B.net2

Presentation Overview

- Introduction & Problem Statement
- Characteristics of IDS alarms
- Filtering False Positives
- Problems with Filtering
- Benefits of Filtering -- Looking at What's Left Over
- Conclusion & Open Problems

Conclusion & Open Problems

- Filtering is very powerful, but comes at a cost
- Based on 7 sensors over 4 months (almost 1 GB):
 - ▶ Removing 60% -- 80% of alarms is possible
 - ▶ at the cost of developing 1 -- 2 new filters per month
- Main problem: Finding filters efficiently
 - ▶ First experiments with visualization and data mining look very promising