Detecting Unknown Network Attacks With Flows

Miguel Pupo Correia
joint work with Luís Sacramento
IFIP WG 10.4 72nd Meeting
Longmont, CO, USA – June 2017

Motivation

- Scenario:
  - Large national telco: mobile commun., Internet, TV,…
  - Connected to its own provider
  - Huge amount of traffic in/out, much is encrypted
  - Possibly new attacks / new variants

2

Motivation

- Compromised hosts do attacks such as:
  - Distributed denial of service attacks
  - Exfiltrating confidential data
  - Sending spam
  - Mapping the network
  - Contact bot command&control centers
- ... and new attacks may appear

3

Network Intrusion Detection Systems

- Traditional NIDSs:
- Knowledge-based: require signatures of attacks
  - Not good for new attacks
- Behavior-based: allow detecting new attacks but require clean traffic for training
  - Where to get it with our scenario?
- Most do deep packet inspection, unfeasible with a lot of traffic

4

Approach

- Framework to detect malicious hosts based on network traffic
  - Not knowledge-based, to avoid need for signatures
  - Not behavior-based, as no training traffic exists
  - No deep packet inspection, as it is slow
- Detects hosts doing new attacks or new variants
- Scale-down data to a level understandable by humans

5

Approach

- Loop:
  - Collect flows for a period of time (e.g., 1 day)

6

Flows

- Flow: sequence of related packets observed during an interval of time
  - A flow is defined in terms of a subset of src IP, dest IP, protocol, src port, dest port; ex: *(*, 1.2.3.4, TCP, *, 80)*
- Netflow: monitoring approach created by Cisco
  - Idea is to capture data about network flows
  - Data: begin/end of flow timestamps, n. packets, n. bytes
  - Variants: IPFIX (standard based on Netflow 9), sFlow,…

7

Approach

- Loop:
  - Collect flows for a period of time (e.g., 1 day)
  - Extract from the flows data about hosts with MapReduce

8

Host data extraction

- Host features (data) extracted by MapReduce:

Feature	Description
Aggregation Key	The IP address that will be used as an identifier, to which the below features relate to
NumSIPs / NumDIPs	The number of IP addresses contacted
NumSports	The number of different source ports contacted
NumDport	The number of different destination ports contacted
textbfNumHTTP	The number of packets to/from port 80 (HTTP)
NumIRC	The number of packets to/from ports 194 or 6667 (IRC)
NumSMTP	The number of packets to/from port 25 (SMTP)
NumSSH	The number of packets to/from port 22 (SSH)
TotalNumPkts	The total number of packets exchanged
PktRate	The ratio of the number of packets sent and its duration
ICMPRate	The ratio of ICMP packets, and total number of packets
SynRate	The ratio of packets with a SYN flag and the total number of packets
TotalNumBytes	The overall sum of bytes
AvgPktSize	The average packet size
BadSubnet	This field expresses whether the IP address belongs to a blacklisted subnet
MaliciousIP	This field expresses whether the IP address is blacklisted
OpenVaultBlacklistedIP	Same as the above, but checked from a trusted and well know threat database
MaliciousASN	This field shows if the IP address belongs to a blacklisted ASN
LocationCode	Code for the country associated with the address

extracted from the flows directly

based on threat intelligence

Approach

- Loop:
  - Collect flows for a period of time (e.g., 1 day)
  - Extract from the flows data about hosts with MapReduce
  - Use clustering to create groups of hosts

10

Unsupervised ML / clustering

- Idea: group similar hosts in clusters (sets)
- Why? Humans can understand and classify a few clusters, not zillions of hosts
- How?
  - Normalize every feature into range [0,1]
  - Run clustering algorithm, e.g., K-Means, to get **k** clusters
  - **k** can be defined, e.g., with the elbow method (finds the "elbow", i.e., when adding more clusters does not improve the modelling of the data)

11

Approach

- Loop:
  - Collect flows for a period of time (e.g., 1 day)
  - Extract from the flows data about hosts with MapReduce
  - Use clustering to create groups of hosts
  - Use classifier to automatically classify hosts
  - Manually label remaining clusters

12

Classification

- End of period 1:
  - we have k clusters => classify them manually
- End of period n:
  - classify hosts or clusters automatically using a classifier (e.g., SVM)
  - classify manually those that do not fit well – human in the loop required as new attacks may exist

13

Approach

- Loop:
  - Collect flows for a period of time (e.g., 1 day)
  - Extract from the flows data about hosts with MapReduce
  - Use clustering to create groups of hosts
  - Use classifier to automatically classify hosts
  - Manually label remaining clusters
  - Repeat for next period

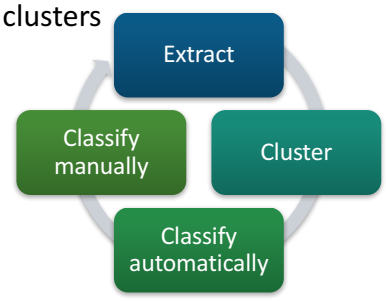

14

Evaluation

- Two parts:
- Synthetic dataset (ISCX)
  - Designed for IDSs
  - Flows are labelled
  - Allows validating the approach
- Real dataset collected at the telco
  - No ground truth

15

ISCX dataset – cluster 3

Clusters	1	2	3	4	5	6	7	8	9	10
# Entries	1	5	1	1	5	5	16	8	3	1
Features										
1	-	-	-	-	-	-	-	-	-	-
2	-	0.125 ; 0.023	-	1.0	0.243 ; 0.064	-	-	0.200 ; 0.016	0.277 ; 0.076	0.115
3	-	-	-	-	-	-	-	-	-	0.227
4	-	0.131 ; 0.023	-	-	0.418 ; 0.56	-	-	0.213 ; 0.031	-	-
5	-	-	-	-	-	-	-	-	-	-
6	-	-	-	0.325	-	-	-	-	-	-
7	-	-	1.0	-	-	-	-	-	-	-
8	-	-	-	-	-	-	-	-	-	-
9	-	-	-	-	-	-	-	-	-	-
10	-	-	-	-	-	-	-	-	-	-
11	-	-	-	-	-	-	-	-	-	-
12	-	-	-	-	-	-	-	-	-	-
13	-	-	-	-	-	0.305	-	-	-	-
14	-	-	-	-	-	-	-	-	-	-
15	-	-	-	-	-	-	-	-	-	-
16	-	-	-	-	-	-	-	-	0.128 ; 0.005	0.184

Avg ; StdDev

Low values

Saturday clusters

- Brute-Force SSH attack found during this day (cluster 3)
  - Maximum for SSH connections (and high, not seen in table)

16

Telco dataset – cluster 21

Cluster #	21
# Hosts	12
Features	
1	-
2	-
3	-
4	-
5	-
6	-
7	-
8	-

Cluster #	21	
# Hosts	12	
Features		
9	0.541 ; 0.181	High IRC communication
10	-	
11	-	
12	-	
13	-	
14	-	
15	-	
16	0.261 ; 0.026	High average packet size

- Bot communicating with C&C server
  - Confirmed by accessing the IP of the C&C server

17

Conclusion

- Network Intrusion Detection for identifying malicious hosts using flows
- ...without having to say how entities misbehave
- Use clustering (unsupervised ML) to reduce the size of the problem and
- a classifier (supervised ML) to automatize classification
- Keep humans in the loop; mandatory w/evolving threats
- Detects attacks involving many packets, not low traffic attacks like buffer overflows or SQL injection

18