

Intrusion Detection -- A 20 year practice

Peng Liu
School of IST
Penn State University

Copyright (c) Peng Liu, The
Pennsylvania State University

1


Outline

- ⌘ Motivation
- ⌘ Intrusion Detection Techniques
- ⌘ Intrusion Detection Products
- ⌘ Some New Research Directions

Copyright (c) Peng Liu, The
Pennsylvania State University

2

Till 1980 ...




Reference monitor

Copyright (c) Peng Liu, The
Pennsylvania State University

3

However attacks do succeed



Copyright (c) Peng Liu, The Pennsylvania State University

4

How serious?

⌘ FBI investigates 500 organizations

☑ 1996, 42% had a security breach

☑ 1997, 50%

☑ 1998, 64%

⌘ Total lost

☑ 1997, 100 million

☑ 1998, 138 million

Copyright (c) Peng Liu, The Pennsylvania State University

5

Next Generation Security

How to detect and respond to these successful attacks such that critical function can be sustained?

Copyright (c) Peng Liu, The Pennsylvania State University

6

What is intrusion detection?

If computer security measures are analogous to the fences and locks of the physical world, then intrusion detection is like a burglar alarm system.

Copyright (c) Peng Liu, The Pennsylvania State University

7


Outline

- ⌘ Motivation
- ⌘ Intrusion Detection (ID) Techniques
- ⌘ Intrusion Detection Products
- ⌘ Some New Research Directions

Copyright (c) Peng Liu, The Pennsylvania State University

8

ID Systems



Copyright (c) Peng Liu, The Pennsylvania State University

9


Example ID Systems

<i>Host based</i>	<i>Network based</i>
ComputerWatch, Discovery, Haystack, IDES (EMERALD), ISOA, MIDAS, USTAT, Wisdom and Sense	ISOA, IDES, NADIR, DIDS, NSM, GrIDS, NetSTAT, NID

Copyright (c) Peng Liu, The Pennsylvania State University

10

Host based ID Systems



Copyright (c) Peng Liu, The Pennsylvania State University

11

Host attacks

- ⌘ Local to Root (L2R)
- ⌘ Attempted break-in
- ⌘ Masquerading
- ⌘ Leakage by legitimate user
- ⌘ Inference by legitimate user
- ⌘ Trojan horse
- ⌘ Virus

Copyright (c) Peng Liu, The Pennsylvania State University

12

Ex 1: Detect masquerading

- ⌘ Masquerading - the attacker uses a legitimate user's account
- ⌘ Observation: the attacker's behavior may differ considerably from that of the legitimate user
- ⌘ The idea
 - ☑ build a **profile** of the legitimate user
 - ☑ if a behavior (on behalf of the user) is very different from the profile, raise a warning
 - ☑ use measures to quantify profiles and behaviors

Copyright (c) Peng Liu, The Pennsylvania State University

13

Host Measures (for a session)

<i>Measure</i>	<i>Description</i>
CPU usage	CPU time
Audit Record	# of audit records (for each hour)
File Usage	# of times each file was accessed
System Errors	# of times each type of error occurred
Directory Usage	Whether a directory was accessed
System Call	# of times each system call was used

Copyright (c) Peng Liu, The Pennsylvania State University

14

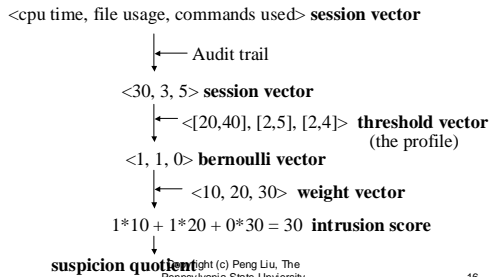
Types of Measures

- ⌘ Intensity measures - the number of audit records per time interval
- ⌘ Audit record distribution measure
- ⌘ Categorical measures - names of files, terminals, and remote hosts
- ⌘ Counting measures

Copyright (c) Peng Liu, The Pennsylvania State University

15

The Haystack Algorithm (1)



16

The Haystack Algorithm (2)

suspicion quotient - the probability that a random session's intrusion score is less than or equal to the session's

Session 1 -- 50
Session 2 -- 30
Session 3 -- 60
Session 4 -- 40
Session 5 -- 10
Session 6 -- 60
Session 7 -- 20
Session 8 -- 30
Session 9 -- 40
Session 10 -- 50


suspicion quotient = 0.40

0 - very suspicious
1 - very normal

Copyright (c) Peng Liu, The Pennsylvania State University

17

The NIDES Algorithm



Copyright (c) Peng Liu, The Pennsylvania State University

18

Ex 1: summary

- ⌘ This approach is an **anomaly detection** approach
- ⌘ This approach is profile based
- ⌘ This approach is a statistical approach
- ⌘ This approach can be used to detect many other kinds of attacks

Copyright (c) Peng Liu, The Pennsylvania State University

19

How good is an ID technique?

- ⌘ Detection rate
- ⌘ False alarm rate
- ⌘ Detection latency

Copyright (c) Peng Liu, The Pennsylvania State University

20

Other Anomaly Detection Techniques

- ⌘ Rule based (i.e., threshold based)
- ⌘ Using neural networks
- ⌘ Temporal sequence learning
- ⌘ Mining profiles

Copyright (c) Peng Liu, The Pennsylvania State University

21

Ex 2: detect a L2R attack

⌘ The attack: (in SunOS 4.1.1)

⊠% ln target -x

⊠% -x

⌘ Note: `target` is a `setuid` shell script owned by the root

⌘ Note: executing `_x` invokes an interactive subshell with root privileges

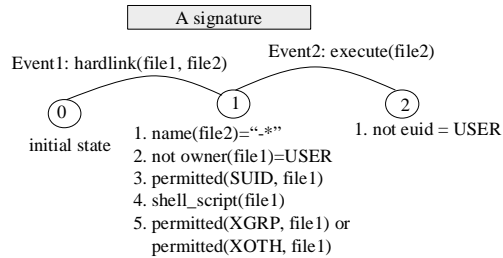
⌘ Observation: all such attacks have the same **pattern**

⌘ The idea: if a behavior matches the pattern, then it is a L2R attack

Copyright (c) Peng Liu, The Pennsylvania State University

22

Using state transition diagrams to specify patterns



Copyright (c) Peng Liu, The Pennsylvania State University

23

Using rules to specify and detect signatures

Event: `hardlink(file1, file2)`

Condition: the resulted state satisfies (1) `name(file2)!="*"`; (2) `not owner(file1)=USER`; (3) `permitted(SUID, file1)`; (4) `shell_script(file1)`; (5) `permitted(XGRP, file1)` or `permitted(XOTH, file1)`

Rule 1

Action: put (USER, file2) into state 1

Event: `execute(file2)`

Condition: (1) the before state satisfies (a) (USER, file2) is in state 1; (2) the after state satisfies (a) `not euid=USER`

Rule 2

Action: raise a warning of **hardlink** attack

Copyright (c) Peng Liu, The Pennsylvania State University

24

Ex 2: summary

- ⌘ This approach is a **signature-based detection** approach
- ⌘ Either rules or state transition diagrams can be used to specify a signature
- ⌘ Detection is done by matching signatures
- ⌘ This approach can be used to detect many other kinds of attacks

Copyright (c) Peng Liu, The Pennsylvania State University

25

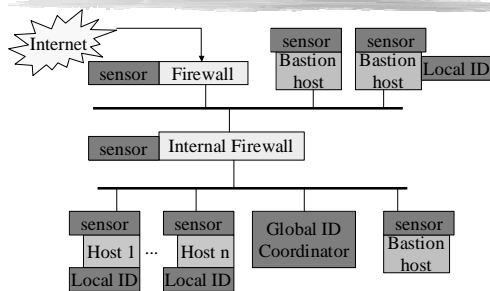
Anomaly detection vs. Signature-based detection

<i>Anomaly detection</i>	<i>Signature-based</i>
Good for unknown attacks	Unable to detect unknown attacks
Limited for known attacks	Good for known attacks

Copyright (c) Peng Liu, The Pennsylvania State University

26

An network-based ID system



Copyright (c) Peng Liu, The Pennsylvania State University

27

Network Attacks

- ⌘ R2L - Remote to Local
- ⌘ U2R - User to Root
- ⌘ DoS - Denial of Service
- ⌘ Probing

Copyright (c) Peng Liu, The Pennsylvania State University

28


Attacks on Solaris Hosts

R2L	Dictionary, Ftp-write, Guest, Phf, Xlock, Xsnoop
U2R	Eject, Ffbconfig, Fdformat, Ps
DoS	Apache2, Back, Mailbomb, Neptune, Ping of Death, Process table, Smurf, Syslogd, UDP Storm
Probing	IP Sweep, Mscan, Nmap, PortswEEP, Saint, Satan

Copyright (c) Peng Liu, The Pennsylvania State University

29

Ex3: Detect TCP Hijacking



Assume A and B are talking via a TCP connection

1. B to A:

A	B	2357	ACK
---	---	------	-----

 An ACK packet
2. Attacker to B:

B	A	2357	data
---	---	------	------

 A request packet
3. B to A:

A	B	2358	ACK
---	---	------	-----

 An ACK packet; A will reject it
4. A to B:

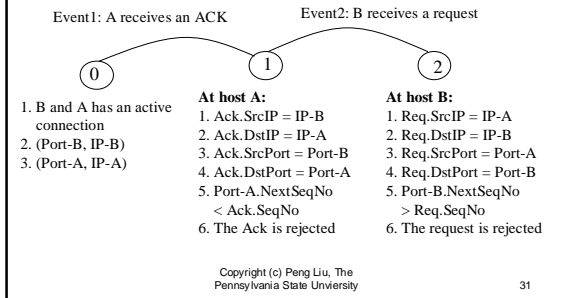
B	A	2357	data
---	---	------	------

 B will reject it

Copyright (c) Peng Liu, The Pennsylvania State University

30

The signature for TCP hijack



Centralized detection

- ⌘ A's (B's) sensor sends each received packet and Port-A's NextSeqNo (Port-B's NextSeqNo) to the ID coordinator
 - ⌘ When the ID coordinator receives the report from A's sensor, it enters **state 1**
 - ⌘ When the ID coordinator receives the report from B's sensor, it enters **state 2** and raises a warning
- Copyright (c) Peng Liu, The Pennsylvania State University 32

Distributed detection

- ⌘ A's sensor does some local ID. When it enters **state 1**, it will inform the ID coordinator
 - ⌘ B's sensor does some local ID. When it enters **state 2**, it will inform the ID coordinator
 - ⌘ When the ID coordinator receives the two reports from A's sensor and B's sensor, it raises a warning
- Copyright (c) Peng Liu, The Pennsylvania State University 33

Centralized vs. Distributed ID

Centralized ID	Distributed ID
	More scalable
	More collaboration overhead
More communication overhead	
More cost	
More accurate	

Copyright (c) Peng Liu, The Pennsylvania State University

34


Ex 3: summary

- ⌘ Network ID needs info from multiple sensors
- ⌘ Each sensor reports two kinds of info
 - ☑ network traffic
 - ☑ host activities
- ⌘ Either centralized or distributed ID are possible
- ⌘ Either anomaly detection or signature-based techniques are useful

Copyright (c) Peng Liu, The Pennsylvania State University

35

ID techniques: A Summary



Copyright (c) Peng Liu, The Pennsylvania State University

36

Outline

- ⌘ Motivation
- ⌘ Intrusion Detection (ID) Techniques
- ⌘ Intrusion Detection Products
- ⌘ Some New Research Directions

Copyright (c) Peng Liu, The Pennsylvania State University

37

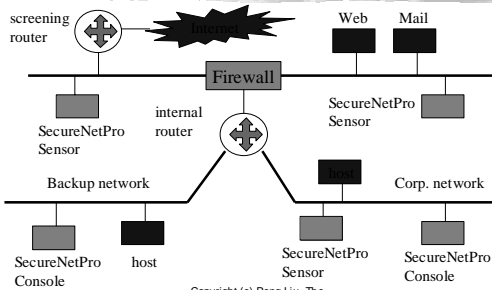
Commercial ID Systems

- ⌘ VCC by Tripwire
- ⌘ CMDS by SAIC
- ⌘ SecureNetPro & Kane Secure Enterprise by intrusion.com
- ⌘ INTOUCH NSA by TTI
- ⌘ NetRanger by Wheelgroup
- ⌘ PolyCenter by Digital
- ⌘ Real Secure by ISS
- ⌘ WatchDog by InfoStream
- ⌘ Stalker

Copyright (c) Peng Liu, The Pennsylvania State University

38

SecureNetPro: An example



Copyright (c) Peng Liu, The Pennsylvania State University

39

SecureNetPro: features

- ⌘ 100 Mbps real-time detection
- ⌘ monitor over 50 segments simultaneously
- ⌘ more than 300 common attack signatures
- ⌘ session replay via TCP/IP reconstruction
- ⌘ stateful application layer protocol decoding
- ⌘ e-mail notification; customizable reports

Copyright (c) Peng Liu, The
Pennsylvania State University

40

Outline

- ⌘ Motivation
- ⌘ Intrusion Detection (ID) Techniques
- ⌘ Intrusion Detection Products
- ⌘ Some New Research Directions

Copyright (c) Peng Liu, The
Pennsylvania State University

41

New Directions

- ⌘ Application-aware intrusion detection
 - ☑ database applications
 - ☑ distributed applications on CORBA or DCOM
- ⌘ Automatic profile and rule management
- ⌘ Collaborative intrusion detection
- ⌘ Advanced ID techniques
- ⌘ QoS of intrusion detection systems

Copyright (c) Peng Liu, The
Pennsylvania State University

42

Ex: Mining rules from trails

- ⌘ Rules are useful for both signature based detection and anomaly detection
- ⌘ Managing rules manually has many limitations: ad-hoc, prone to errors, etc.
- ⌘ The idea
 - ☑ mine signatures for known attacks from network traffics and host trails
 - ☑ mine profiles from legitimate network traffics and host trails

Copyright (c) Peng Liu, The Pennsylvania State University

43

Acknowledgements

- ⌘ Some materials of this lecture are from Prof. Peng Ning at North Carolina State University

Copyright (c) Peng Liu, The Pennsylvania State University

44
