15

PANEL ON INTRUSION DETECTION

TC Ting, Ming-Yuh Huang, D. Shayne Pitcock, Chris Clifton
and T.Y. Lin, Chair

Abstract We try, in this panel, to give an overview of intrusion detection, and then propose some areas where work in database security and intrusion detection can and should overlap.

Keywords: Data analysis, intrusion detection, operation model, vulnerability management

1. T. C. TING

What is Intrusion?

To detect an intrusion, we must know what intrusion is. Intrusion can be defined as entry to disrupt or gain illicit access to a systems. It is one form of "infowar."

An intrusion can be caused by

1. External perpetrator
   - Acting as seemingly authorized users.
   - Engaging in passive and/or active wiretapping.
   - Sending virus/worms, agents, and other malicious attacks.
2. Authorized users:
   - Exceed the legitimate privileges. We need stronger access control
   - Inference protected information. We need better statistical and semantical inference detection and control.
3. ACCIDENTAL/INCCIDENTAL ENTRY INTO THE SYSTEM.
   - Today's Internet environment, these type of entries is very frequent.

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: 10.1007/978-0-387-35508-5_22

V. Atluri et al. (eds.), *Research Advances in Database and Information Systems Security*

One way to guard against intrusions is to make authorization/authentication an on-going process in a run-time environment. This can be supported in various ways:

- By message. Public key infrastructure and digital signature are examples.
- Re-authentication. Conducted periodically or when in doubt.
- Intruder detection. Monitoring user and program behavior.

Why is intrusion detection important? After all, once an intrusion has happened, isn't it too late? This attitude doesn't take into account the real need for damage *control* and damage *recovery.*

- Damage Control. It is important to have effective damage control. It is difficult since an intrusion may have

    Long detection latency

    Some intrusion may introduce delayed actions

    Many potential false detections are possible.

- Damage Recovery. It is important and difficult to recover damage since it will be difficult to know what has been damaged. Two potential improvements are: (1)Better back up system and (2)Quick restart

It is important that we keep in mind what we will be doing when an intrusion is detected; this shapes the requirements of an intrusion detection system.

At the final point of the discussion, I raised the issue that the weakest link is the system's administrators who have potential to damage the system but are not legally liable.

2. SHAYNE PITCOCK

What should a complete intrusion detection system (IDS) entail?

2.1 MOTIVATION

Today's network and operating system intrusion detection systems do not adequately address the totality of threats to enterprise-wide computing infrastructures. Traditionally, IDS focused either on monitoring the network traffic or operating system activities. Intrusion attempts that occur at the network level may not be considered an attack at the system level and produce a false positive alert. The opposite occurs at the operating system level: OS attacks may not be considered network intrusions and likewise produce a false positive. Realistically, correlation of multiple data points are required for the most complete and accurate reporting of security alerts.

The success of correlating IDS alerts relies heavily upon knowledge of various data points. This could include data points such as "normal" traffic on a particular network segment, type of data transferring to a particular system, processes allowed on that system, company security policies or the vulnerabilities of the network/system under attack. Each presents a different focus depending upon the monitoring approaches such as external perimeter systems or specific internal business assets.

The proposed approach is for an IDS to capture enough data points for more accurate intrusion alerts by incorporating the best tools from multiple suppliers. Various early IDS implementors who took a network-first approach now recommend a more balanced approach with both network and server based intrusion detection capabilities. Although the technology is too immature to provide comprehensive data correlation, the opportunity does exist to significantly enhance the current enterprise-wide ID capabilities.

2.2 APPROACH

Working scopes. Establish a security solution that provides continuous, near real-time, intrusion detection and response capabilities. The solution should create an "early warning system" to provide a basis for appropriate responses against computing infrastructure attacks. Indirectly, the system would also provide a secondary metrics measurement for the integrity of the company's network and servers.

Architecture. The associated items of a system description (see Figure 15.1) are: Electronic sensing, electronic triage, correlation triage, forensics analysis, computing response, incident response, technical industry monitoring and improvement, and vulnerability management. Interrelated functional components operate independently. Initial identification of computing attacks occurs with the software agents that perform the electronic sensing. The "raw" alerts, such as system configuration deviations, are transmitted to the management consoles that provide the electronic triage and multiple alerts collection. When thresholds are exceeded, or the severity of alert is critical, the alarms become "filtered" and are forwarded to the Network Management Center (NMC) for more detailed correlation triage. The correlation triage begins to consider the impact of security alerts with the possible network alerts reported to the NMC. At this point, depending on the severity and requirement for more detailed response or investigation, the data is forwarded as an incident alert. This is verified by the forensics analysis team and acted upon by the computing response team.

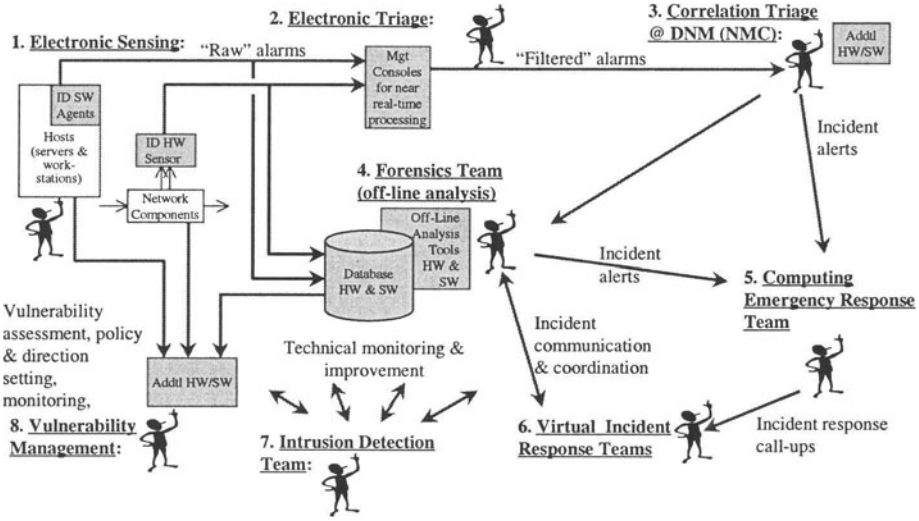
1. Electronic Sensing:
2. Electronic Triage:
"Raw" alarms
Mgt Consoles for near real-time processing
"Filtered" alarms
3. Correlation Triage @ DNM (NMC):
Addtl HW/SW
ID SW Agents
Hosts (servers & work-stations)
ID HW Sensor
Network Components
Incident alerts
4. Forensics Team (off-line analysis)
Off-Line Analysis Tools HW & SW
Database HW & SW
Incident alerts
5. Computing Emergency Response Team
Vulnerability assessment, policy & direction setting, monitoring,
8. Vulnerability Management:
Addtl HW/SW
Technical monitoring & improvement
Incident communication & coordination
7. Intrusion Detection Team:
6. Virtual Incident Response Teams
Incident response call-ups

Figure 15.1 Intrusion detection system.

An incident may require more expertise from particular subject matter experts, as well as a reply from the legal or public relations offices, creating a virtual incident response team. Incidents may require additional expertise/expert on particular subject matter areas or legal/public relations officers in a virtual incident response team. A technical team, which is knowledgeable of IDS industry advancements, also provides consultation for technical improvement opportunities along the process. For comprehensive system coverage, vulnerability assessment and risk management must also be included to integrate security policy, company security direction, and system awareness. All components contribute to better real-time security management on today's distributed computing infrastructures.

3. CHRIS CLIFTON

Where does *database technology* come into this?

First, as Shayne showed us in Figure 15.1, there is a lot of data associated with an IDS. Storage and manipulation of this information is a database problem. Two database research areas are particularly affected:

Active Database Quick, automated response is necessary. Although some response may be based on "raw" alarms, more sophisticated responses may be based on data from multiple sources. Active database technology can provide an efficient, effective means for implementing such response.

Real-time database Automated response is most needed where timely action is required. Databases that do not support real-time requirements will be relegated to supporting off-line analysis.

Solving problems in these areas will allow better support for steps 2 and 3 in Figure 15.1.

Database technology can serve in other ways as well. One point often lost is that **intrusion $\neq$ compromise of data**. Databases provide additional security features. The key is to

1. Authenticate database use separately from the operating system, and

2. Don't give users more authorization than they need.

In practice, we fail on item 2. This is a research challenge: existing authorization mechanisms are too difficult to administer. This conference presents improvements (e.g., the session on role-based access control), but more work is needed. We need to ensure that in the real world, *we close the barn doors leading to the data.*

Even more advanced is to apply intrusion detection to the database itself.[1] By detecting abnormal use of the database, we can potentially catch insiders misusing the data as well as intruders.

Another area where the database community can help is with off-line analysis of intrusion detection data. Data mining technology can help us build better IDS filters, *if we know how to apply it.* As one example, we are using sequential association mining at MITRE to identify frequently recurring patterns of alarms. The goal is to identify false alarms – enabling us to build better filters. The key is that we don't expect the data mining tools to identify false alarms. Instead, we expect them to identify the best places to apply human expertise. By manually analyzing the *most common* patterns, we can be assured that the reward (in reduction of false alarms) will be highest for the patterns we *can* identify as representing normal behavior. There are many other potential uses of data mining to support intrusion detection, for an example see [2].

4. T. Y. LIN

Data Mining Research for Intrusion Detection.

Data mining provides substantial capabilities to identify patterns in data. This can have substantial applicability to intrusion detection. However, some of the problems in intrusion detection (and some medical applications as well) are somewhat more delicate. We need a clearer notion on *what is a pattern?* Some intuition may be deceiving: Suppose we are given a sequence [3].

$$1,3,5,7$$

What is the next number? To get the answer, one would find a pattern, namely, a formula that generates the sequence. ¿From MATLAB, however, we find the following "surprises:"

1. $f_6(n)$ predicts: 1,3,5,7,6.

   $f_6(x) = -0.1250x^4 + 1.2500X^3 - 4.3750X^2 + 8.2500X - 4.0000.$

2. f_{15} predicts: 1,3,5,7,15

   $f_15(x) = 0.2500x^4 - 2.5000X^3 + 8.7500X^2 - 10.5000X + 5.0000.$

Given a set of data, there is a complexity associated to it, called data complexity. A pattern supported by this set has its own complexity, called pattern complexity. If the pattern complexity is strictly smaller than data complexity, we will say the pattern is a pattern. Taking this view, we can conclude that those "surprises" are not patterns. In low frequency data, patterns can easily be fitted erroneously. There has been work in applying data mining to intrusion detection [2]; this is a good area for further work.

5. AUDIENCE

Joachim Biskup. What is "unauthorized access"?

- Ming-Yuh. Defining access requires a policy, a trust model, and identity (X.509).
- Shayne. There are multiple points of collection for these – need to know what information is needed.

Reind van de Riet. How many security people are there at Boeing?

- Ming-Yuh. Can't answer that question.
- Shayne. Security is recognized as important (I report to the CIO). *Everybody is responsible for security.*

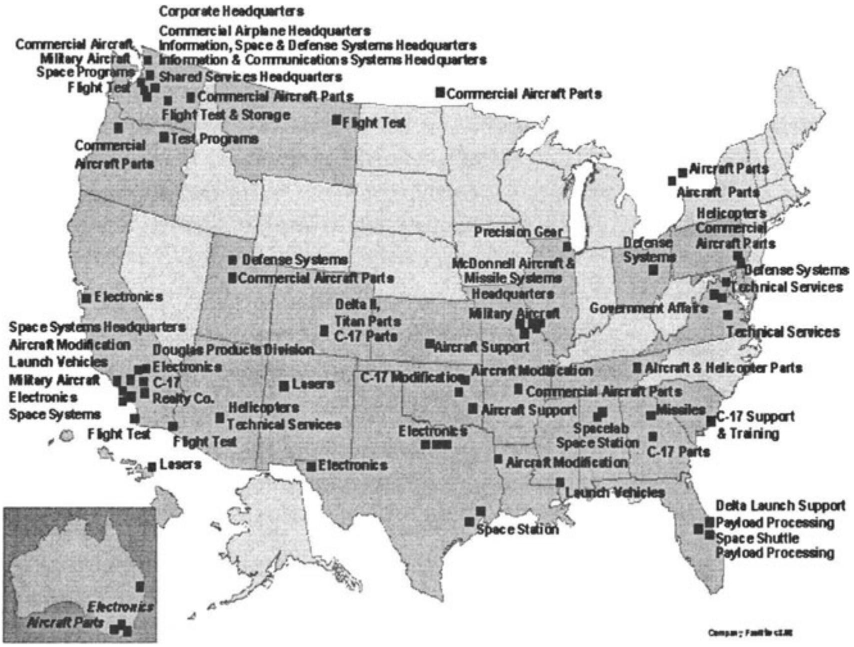

Figure 15.2 Boeing campuses.

Ravi Sandhu. Do you use encryption at Boeing?

- Ming-Yuh. Encryption is used between campuses. (Note that this is a big issue, see Figure 15.2.)

T.C. Ting. Who do we trust?

- Chris. Mechanisms exist to split trust (e.g., key escrow) – we don't have to put all our trust in any single entity.
- John Dobson. *Social processes* of trust are not well understood.

References

[1] Chung, Gertz, and Levitt (1999). Misuse detection in database systems through user profiling. *Proceedings of the Second International Workshop on Recent Advances in Intrusion Detection.*

[2] Lee, W., Stolfo, S. J., and Mok, K. W. (1998). Mining audit data to build intrusion detection models. *Proceedings of the Fourth International*

Conference on Knowledge Discovery and Data Mining (KDD-98) (eds. R. Agrawal, P. E. Stolorz and G. Piatetsky-Shapiro), pp. 66–72.

[3] Lin, T.Y. (1999). Discovering patterns in numerical sequences using rough set theory. *Proceedings of the Third World Multi-conference on Systemics, Cybernetics and Informatics.*