Intrusion Detection for an On-Going Attack

Jim Yuill
jimyuill@pobox.com
Computer Science Dept.
North Carolina State Univ.

S. Felix Wu
wu@adm.csc.ncsu.edu
Computer Science Dept.
North Carolina State Univ.

Fengmin Gong
gong@mcnc.org
Adv. Networking Research
MCNC

Ming-Yuh Huang
ming-yuh.huang@boeing.com
Applied Research and Tech.
The Boeing Company

Abstract

An intrusion-detection system (IDS) for an on-going attack is described. Prior to an attack, an IDS operates in anticipation of a general threat. During an attack, the IDS can deal less in the general and more in the particular- namely, particulars about attackers and attacked devices. A profile of the attacker is developed, using information he reveals about himself during his attacks. Principles from economics are used to predict the attacker's behavior, based on estimates of his asset-appraisal, attack-costs and attack-resources. Likely-compromised devices (LCDs) are identified, using the profile and the economics-based estimates. Knowledge of LCDs is useful for work in attack repair, neutralization and containment.

I. INTRODUCTION

When a network is under attack, an intrusion-detection system faces unique difficulties and opportunities. This paper explores those difficulties and opportunities, and it presents an intrusion detection technique based upon them. The technique is descriptively named ***Investigative Intrusion-Detection (I-ID)***.

The difficulties

A sys-admin discovers that a device on his network has been compromised. To repair that device, he performs, roughly, these tasks: 1) remove the attacker's active processes, 2) assess and repair damage from the attack, 3) determine the compromised vulnerability, 4) use risk analysis to choose appropriate countermeasures for the vulnerability, 5) remove the compromised vulnerability, by repairing or improving the system.

However, after securing the compromised device, the sys-admin himself does not rest securely. One compromised device on a network raises questions about the security of the entire network. For the attackers who have compromised this device, what is their other harmful activity on the network: past, present and future? Having discovered one compromised device, the sys-admin must ensure the rest of the network is secure. Furthermore, when an unsuccessful attack is discovered, the sys-admin can have similar concerns.

Here's an example of what a sys-admin fears when he discovers a compromised device. In 1998 a large telephony OEM discovered that one of its networks was compromised. The attacker had extensively compromised the network over a long period, prior to being discovered. Repairing the network and removing the attacker was a costly and lengthy process.

The opportunities

A network is attacked by a particular set of individuals. During the attack, each individual reveals information about himself. This information can be used as a means for intrusion detection.

Similarly, particular devices are attacked on the network. Information about these devices can also be used as a means for intrusion detection.

Prior to an attack, an intrusion detection system operates in anticipation of a general threat. During the attack, the intrusion detection system can deal less in the general and more in the particular- namely, particulars about attackers and attacked devices.

Here's an example of how ID information can be obtained from an attacker's activities and from an attacked-device. In 1999 a research organization discovered that two of its Linux machines were compromised. Investigation revealed that a single attacker had compromised both boxes, using the same exploit on each machine. The organization then investigated its other Linux machines to determine if they had the same vulnerability and if they were compromised. The attacker's telnet sessions were put under surveillance. The attacker was observed running "bots" to attack IRC servers. He was not observed attacking other machines in the organization. It was concluded that the attacker was probably a "script kiddy" and that he had no particular interest in this organization.

The objective

Three measures used to secure a compromised network are:

?? ***attack repair***: repairing devices altered by the attacker

?? ***attack neutralization***: fixing vulnerabilities which the attacker has exploited, or which he will exploit

?? ***attack containment***: temporary measures for limiting an active attack, e.g. blocking all ftp sessions at the firewall

We'll refer to "attack repair, neutralization and containment" as ***ARNC***.

ARNC can be a difficult task for the sys-admin, especially in a large heterogeneous network. Our intrusion detection technique intends to help the sys-admin with these problems:

?? The sys-admin needs to find the devices that have been, or will be, compromised in the attack.

?? When under attack, speed is an important element of ARNC tactics. The sys-admin needs to find and fix vulnerabilities before the attacker can exploit them.

The scope of our consideration is an attack against a single domain. The attack is between network-attached devices. We are not considering social-engineering attacks. The attack can originate inside or outside the domain. It is assumed that there is at least one known, or suspected, attack.

II. AN OVERVIEW OF THE INVESTIGATIVE INTRUSION-DETECTION TECHNIQUE

In this section, an overview of the intrusion detection (ID) technique is presented. The next section elaborates on its parts. The technique is descriptively named ***Investigative Intrusion-Detection (I-ID)***.

Likely-Compromised Devices (LCD's)

For ARNC, the sys-admin would like to find all the devices that have been compromised. However, investigating devices for compromise can be a labor-intensive and lengthy task: Devices can be checked manually for telltale signs of compromise, e.g. strange accounts in /etc/passwd. ID systems (IDS's) can be configured to be more sensitive or to look for specific indications of compromise. IDS's that run periodically, e.g. tripwire, can be run immediately.

It is difficult, perhaps impossible, for the sys-admin to investigate every network device for compromise. To assist the sys-admin, we would like to identify the network devices that are likely to have been compromised. Also, we would like to identify the degree of likelihood. Prioritizing ARNC-work according to the most-likely-compromised devices allows us to make effective use of the sys-admin's limited resources for investigation. Prioritization also improves the sys-admin's speed, an important tactic in ARNC.

We'll refer to ***likely-compromised devices*** as ***LCD***'s.

For ARNC, the sys-admin would like to know which devices would be compromised in the future. Here too, speed is important-- the sys-admin would like to fix vulnerabilities before they are compromised. We'll help the sys-admin predict devices that are likely to be compromised in the future, and their degree of likelihood.

The objective of the I-ID technique is to identify LCD's, and their degree of likelihood. Collectively, these LCD's form the ***LCD-set***. A device can be considered likely-compromised for more than one reason. The primary reasons are a vulnerability to attack (e.g. buffer-overflow) or something on the device which the attacker values (e.g. credit-card database). Elements in the LCD-set will be uniquely identified by two attributes: 1) the device, and 2) the reason for suspecting compromise. We'll elaborate upon these reasons later.

Building the LCD-set

The I-ID technique consists of two parts, ***evidence-collection*** and ***LCD-prediction***. In evidence-collection, we gather information about attacks and attackers. The attacks, and attackers, can be known, or suspected. In LCD-prediction, the evidence is analyzed to identify LCD's.

Evidence-collection provides a framework for obtaining the information needed for LCD-prediction. The framework identifies that information, provides a process for obtaining it, and provides a means for storing it. Evidence-collection consists of two parts: ***attacker-profiling*** and ***economic-attribute appraisal***.

During an attack, the attacker reveals information about himself and the attack. We'll use that information to build a profile of the attacker. The profile consists of attacker-attributes that are useful for identifying LCD's. For example, the attacker may reveal the exploits which he uses, his skill level, his objectives, etc.

Economic-attribute appraisal uses principles from economics to understand and predict the attacker's behavior. The attacker's economic-attributes are: 1) his valuation of network assets, 2) his costs for exploiting them, and 3) the attack resources he possesses. The profile, and information about the network, will be used to estimate the attacker's valuation of assets and his exploitation costs.

Devices with known vulnerabilities have finite attack costs. Part of economic-attribute appraisal is ***vulnerable-device detection***. To identify these devices, we'll use known attack techniques combined with information about the devices that are known, or suspected, to be compromised.

In LCD-prediction, the collected evidence is used to identify LCD's and their degree of likelihood. Some principles and heuristics are provided for prediction. The prediction will also depend on the sys-admin's subjective assessment, but it will be informed by the collected evidence.

An environmental constraint: uncertainty

The collected evidence is analyzed to understand and anticipate the attacker's behavior, and thus identify likely-compromised devices. The analysis is speculative rather than deterministic. It is expected that the analysis will sometimes be incorrect. However, to be useful, the analysis need only be better than the alternative techniques (i.e. ad hoc) and be worth the cost of performing it.

We believe that uncertainty is inherent in the prediction of future human behavior. Also, the prediction of past behavior is likely to be uncertain when working with incomplete information. These sources of uncertainty necessitate the use of speculative means for evidence-collection and LCD-prediction. In military combat and in entrepreneurial ventures, success depends upon wise speculation about future human behavior. We believe ARNC and Investigative Intrusion-Detection are similarly speculative.

III. COMPONENTS OF THE I-ID TECHNIQUE

This section describes the components of the Investigative Intrusion-Detection (I-ID) technique. In overview, the components are:

?? ***evidence-collection***: information is gathered about attacks and attackers. It consists of two parts:

- ?? ***attacker-profiling***: a description is made of the attacker, which is useful for identifying LCDs. Due to the length of the description, it is presented in three sections: 1) ***an introduction***, 2***) profile-attributes***, and 3) ***multiple-attackers***
- ?? ***economic-attribute appraisal***: principles from economics are used to understand and predict the attacker's behavior. It contains one sub-part:
    - ?? ***vulnerable-device detection***: vulnerable devices are found using known attack techniques, combined with information about devices which are known, or suspected, to be compromised

?? ***LCD-prediction***: the collected evidence is used to identify LCD's and their degree of likelihood

Attacker-profiling: an introduction

As part of the I-ID technique, attacker-profiling serves two purposes. In economic-attribute appraisal, it provides a means for accurately estimating the attacker's valuation of network assets and for estimating his costs and resources. In LCD-prediction, attacker-profiling is used as a means for anticipating likely attacks.

Generic and customized profiles

Two types of profiles are used, ***generic*** and ***customized***. A generic profile is constructed for typical groups of attackers. Presently, we plan to develop profiles for crackers, criminals and vandals. [ICOVE] Other profiles are possible, e.g. "script kiddies".

In attacking, an attacker necessarily reveals information about himself. Using this information, we'll construct a ***customized*** profile for a particular attacker. Other sources of information may also be available, e.g. a priori knowledge of an attacker who is a known insider.

A single ***profile-template*** will be created for use with both generic and customized profiles. A customized profile can be started by "inheriting" the attributes of a generic profile. The attributes can then be updated as more is learned about the attacker.

Profile organization

First, we'll develop the profile for a single attacker. In a later section we'll develop the profile for a group of attackers.

The profile is made of ***profile-attributes.*** They are the attacker characteristics which, ultimately, are used for identifying LCDs. As mentioned, the attributes are used for economic-attribute analysis and LCD-prediction. We also need profile-attributes for identifying individual attackers, when multiple attackers are present.

There are countless attributes that can potentially be used for identifying LCD's. Our intention is to just identify the most common and useful profile-attributes. To accommodate the remaining attributes, there should be a mechanism for user-defined attributes.

The profile-attributes are organized in categories which correspond to characteristics of the attacker himself: 1) what he has done, 2) what he can do, 3) what he does, 4) what he knows, 5) what he wants, and 6) what identifies him. Some of the categories contain sub-categories. We were unable to find an ideal taxonomy-- some of the categories and sub-categories overlap, e.g. "what he can do" can be a means of identification, overlapping "what identifies him".

These attribute-categories are developed in the next section.

Attacker-profiling: profile attributes

Below, there is a section for each category of profile-attributes. The section-heading names, and describes, the category.

History: what the attacker has done

The profile is based on historical information about known, or suspected, attacks and attackers. For the purpose of building the profile, any potentially useful information should be kept, provided it's cost-justified. It can be obtained from the established sources for intrusion information: log records, surveillance of network connections, manual investigation of compromised devices, IDS's, etc.

Some of the information to be recorded is:

?? information about known attacks, both successful and unsuccessful

?? information about an identified-attacker's activity, both authorized and unauthorized

?? information about known-compromised devices (KCDs). An ancillary purpose of this information is to distinguish KCDs from LCDs.

?? information about devices which are suspected to be compromised

?? a record of ARNC work. This documents the interaction between attacker and defender. It also records when a device is no longer compromised.

Attack evidence varies from being certain to being highly speculative. The certainty of the evidence will be recorded. That information is needed when reasoning from the evidence. Also, the evidence can be incorrect or incomplete. A mechanism is needed for accommodating corrections to the evidence. The accommodation of uncertainty and correction is something that's needed throughout the I-ID technique.

Abilities: what the attacker can do

This section has to do with the attacker's ***abilities***. Knowledge of his abilities is used in economic-attribute assessment for estimating both attack-costs and the attacker's resources. The profile attributes are:

a) The attacker's ***computer-skill*** is a function of his general technical abilities. For example, the attacker may demonstrate strong Unix skills and weak Windows skills. He may have the skill needed to write buffer-overflow attacks, or be scarcely able to run downloaded scripts.

b) ***Attack-skill*** is a function of the attacker's ability to find and exploit vulnerabilities. For example, the attacker may demonstrate knowledge of a wide array of exploits. He may show himself to be very clever and resourceful, or he may work in a very route manner.

c) The attacker's ***tenacity*** describes his persistence. A good example of tenacity is the hacker Matt Singer, from the book *At Large*. [FREED] Although he had low skill, his great tenacity enabled him to penetrate countless systems.

d) The attacker's ***discipline*** describes his organizational skills and thoroughness. Matt Singer was unorganized. He would have been unable to carry out an attack that required meticulous record-keeping.

Computer-skill and attack-skill can be augmented by the attacker's advisors. For example, when Matt Singer was a neophyte, he was assisted by the advice of a skilled attacker. An advisor can be suspected when an unskilled attacker performs a feat that is beyond his normal abilities.

MO: what the attacker does

The profile-attributes for this category describe the attacker's ***MO***, or ***method of operation***. In the field of Criminal Investigation, the ***MO*** is closely related to profiling. It is a summary of the habits, techniques and peculiarities of the criminal. The MO is used to identify a criminal and to predict his behavior. [OHARA]

The scope of consideration is the attacker's tactical behavior. That is, his behavior during an attack. We'll consider his strategic behavior in the profile category "**Motives**: what the attacker wants". That is, his behavior related to his overall goals and objectives.

Some of the profile-attributes for the MO are:

?? exploits used, e.g. a particular buffer-overflow
?? tools used, e.g. the nmap port-scanner
?? techniques for avoiding detection, e.g. erase log-file entries
?? the degree of caution exercised, in avoidance of detection
?? attack technique, e.g. a port-scan followed immediately by an exploit
?? time spent on the network, both duration and patterns-of-occurrence, e.g. time of day. This will help predict the rate at which he works, the volume of his work, and when he is present.

Knowledge: what the attacker knows

The attacker's ***knowledge of the network*** will limit and influence his decision making.

Some types of knowledge included in the profile are:

?? knowledge of the ***network topology***. This includes all network system-information, such as end-hosts and the network infrastructure (e.g. routers). It also includes knowledge of vulnerabilities. For example, we may know the attacker has performed a network-wide scan and identified several devices running a particular Trojan Horse.
?? knowledge of the ***network contents***. This includes all the data stored and processed in the network. For example, an insider may have extensive knowledge of what information is stored in the network, and where.

When recording knowledge, both true knowledge and misperceptions should be recorded. For example, an attacker is compromising devices in the accounting-department's subnet. He thinks credit-card numbers are stored there, but they are actually stored elsewhere.

When recording knowledge, understanding should also be considered. For example, a pharmaceutical company's research subnet is under attack. However, the attacker is a high school student who does not understand the research data nor its value. His lack of understanding will influence his assessment of assets on that subnet.

Motives: what the attacker wants

In general, it is easier to profile behavior than it is to profile psychological attributes like knowledge, motive and personality. Behavior can be objectively observed. Psychological attributes tend to be known through subjective speculation.

There are profile-attributes related to what the attacker wishes to acquire:

?? ***motive*** is the primary distinction used in the traditional categorization of crackers, vandals, and thieves [ICOVE].

?? ***goals*** are the specific objects sought in the fulfillment of motive. A useful distinction in the attacker's goals is whether he is attacking this network in particular, or whether he chose it capriciously. If he chose it capriciously, it's less likely that he would expend a lot of resources to obtain an asset that is available elsewhere at lower cost.

?? ***plans*** on a larger scale may be apparent from actions taken on a smaller scale. For example, an attacker may compromise a device that clearly is of no immediate value to him. It may then be suspected that he is seeking to use the compromised device as a means for attacking some other device.

There are profile-attributes related to what the attacker does not wish to lose:

?? ***personal-assets*** are things subject to loss if the attacker is caught, e.g. employment, freedom from incarceration, and reputation.

?? ***attack-assets*** are things the attacker has acquired on the victim network. For example, the attacker may have invested much time to compromise a device from which he can sniff a valuable network segment.

Attributes about the attacker's ***personality*** can also be useful:

?? ***age*** and ***maturity*** influence motives and goals

?? the attacker's ***culture*** may influence his attack behavior. For example, it may be that attackers from Tokyo and New York City differ in their degree of disregard for their victim.

Psycholinguistics is a tool from the field of criminal investigation. It seeks to understand a person by analyzing his writing. As a simple example, diction indicates what country or region a person is from.

Identifying traits: what identifies the attacker

As described in the next section, it is not uncommon to encounter multiple attackers. To build profiles for individual attackers, we'll need to uniquely identify them. Many of the profile-attributes described earlier can also be used for identification. Some additional attributes are:

Peculiar work-habits can uniquely identify attackers**.** For example, the attacker may habitually use an esoteric option for a particular shell command. An attacker may repeatedly display some illogical behavior. For example, erasing the entire file system, except for the /tmp directory.

The overall objective of the I-ID technique is to identify LCDs. Once identified, the LCDs will need to be investigated for compromise. Investigative techniques are beyond the scope of this paper. Profiling, especially for attacker-identification, could be extended for use in investigation. For example, if an attacker always creates a directory in /tmp named '.. ' (dot, dot, space), then that name could be used to identify him, for the purposes of both profiling and investigation.

Attacker-profiling: multiple attackers

Reports from major hacking cases reveal that hackers often work in small and loosely-formed groups. Collaboration is almost essential for acquiring a high degree of hacking skill. Collaboration is also a means for accomplishing feats that are beyond the ability of an individual.

We would like to identify individual attackers, groups of attackers, and the nature of collaboration within a group.

The relationship between attackers and devices is many-to-many. A device can be attacked by a single attacker or multiple attackers. A single attacker can attack one, or many, devices.

Cooperation among attackers varies. There may be no cooperation; the presence of multiple attackers can be coincidental. If there is cooperation, it can vary from casual acquaintances to well-organized gangs.

When multiple attackers cooperate, a profile of the group can be a means for intrusion detection. Ultimately, groups are made of individuals. Individual profiles (described earlier) will be the basis for the group profile.

Loosely organized groups can be profiled by recording each attacker's ***known accomplices***. This information can be kept in the individual profile. We may assume that accomplices share knowledge of exploits, vulnerabilities and assets. Accomplices differ from advisors (mentioned earlier) in that an accomplice attacks with the attacker to achieve a common goal, and an advisor just provides information. One person can be an attacker's accomplice and his advisor.

We will seek to construct a profile of a group with a common goal. The attributes of a ***group profile*** are:

1. the group's membership:
    a) For each identified member of the group there should be an individual attacker profile.
    b) Leaders of the group should be identified.
2. an estimate of the group leaders' motives, and their goals and plans for the group

Economic-attribute appraisal

The Economics of Crime uses principles from economics as a means for understanding and predicting criminal behavior. For our purposes, the economic attributes relevant to an attacker are: 1) his ***valuation of network assets***, 2) his ***costs*** for attacks, and 3) his ***resources*** for attacks. Using information from attacks, we'll attempt to determine those attributes of the attacker.

For each device on the network, we'll ***appraise*** its value to the attacker. A device can have more than one asset, so each asset will be appraised separately. Appraisement can be done prior to the attack using profiles of typical attackers (e.g. crackers, criminals and vandals [ICOVE]). During an attack, we'll attempt to build a more accurate profile of each attacker, in order to provide a more accurate appraisement.

For example, consider a network that contains an ingenious client-server system and a simple web server. For a highly skilled cracker, we might anticipate the client-server system to be of high interest and the web server to be of low interest. For a low-skilled vandal we might anticipate the opposite.

Attack costs include such things as: necessary skill-level, time required, and the attacker's likelihood of detection. Using our knowledge of devices that are known, or suspected, to be compromised, we'll identify other vulnerable devices. Vulnerable devices are identified using known attack techniques. This is described in the next section, "Vulnerable-device detection". We'll estimate the attacker's cost for each vulnerability discovered on a device. For devices with no known vulnerabilities, we'll consider the attack cost to be infinite.

For example, consider a device X which is the known victim of a remote root-shell exploit. Using the showmount command, it is discovered that device Y allows X to remotely mount its C:\ directory. Y would be considered a vulnerable device. The cost of attack would be low.

The third component is the attacker's resources. Using information from attacks, we'll estimate the attacker's resources. Resources include such things as: the amount of time available, abilities, knowledge of the network, shell accounts, and accomplices.

For example, consider an attacker whose activity has been recorded via log records and surveillance of telnet connections. They reveal that he is present one to three hours per week, and that he is very familiar with the network topology. Also, he primarily accesses Linux devices and rarely accesses Windows devices. These attributes provide insight into the type of attacks that he likely can, and will, perform.

The attacker is not omniscient.   His knowledge of the network can be incomplete or incorrect. This will affect his valuation of assets and estimates of costs. Our analysis of the attacker's economic-attributes should anticipate his ***imperfect knowledge***.

The number of incorrect notions about a thing far exceeds the number of correct notions. Analysis of economic attributes should, in most cases, first assume the attacker has correct and complete knowledge. From this analysis, likely incomplete and incorrect knowledge can be anticipated.

For example, when estimating costs, we'll first try to identify all vulnerabilities. We can then estimate the attacker's likelihood of discovering them. When appraising assets, we'll first assume the attacker knows they exist. We can then estimate his likelihood of discovering them.

Vulnerable-device detection

As mentioned earlier, to estimate attack costs we need to identify vulnerable devices. One way to find vulnerable devices is to use a vulnerability scanner on the network, e.g. ISS' Internet Scanner (TM).

Another way is to use the set of devices that are known, or suspected, to be compromised. We'll refer to this as the ***CD*** (Compromised Device) set. Using known attack techniques, we can identify devices that are vulnerable to attack from a CD, and devices from which the CD is vulnerable to attack.

System-software and configuration errors

A common means of compromise is system-software errors or configuration errors. When the CD is known to be compromised via one of these vulnerabilities, we can search the network for other devices with the same vulnerability.

The sys-admin may not be able to discover the vulnerability that the attacker exploited. It would be reasonable, then, to be suspicious of devices with similar O/S, network services, or configuration.

Vulnerable trust-relationships

Networks are built to share data and services. Security measures hinder sharing. Insider threats are usually lower than outsider threats, so networks often have less security on the inside than on the outside. This configuration is humorously termed, "crunchy on the outside, chewy on the inside."

To better share data, a device may intentionally permit itself to be vulnerable to other devices, especially those inside the network. Some examples are: NFS permissions, r-command permissions, anonymous ftp, and "guest" login-accounts.

Vulnerable trust relationships can be used to identify:

1. devices which may have been attacked from the CD
2. devices from which the attacker accessed the CD

Exploitable information

The CD may contain information whose disclosure renders other devices vulnerable. Here are some examples:

1. A sniffable network connection can provide sensitive information, e.g. telnet user-ids and passwords.
2. Mail folders or text files may contain user-ids and passwords for other devices.

3. If the "crack" program can reveal a password on the CD, other devices that use the same password will be vulnerable. It is common for users to have the same user-id and password on multiple devices.
4. An attacker can exploit information that identifies network assets and vulnerabilities. Examples include a sys-admin's network-topology map and the output of a vulnerability scanner.

Surveillance of the CD

The CD can be put under surveillance, e.g. by sniffing telnet sessions. The attacker's traffic to, and from, the CD may reveal vulnerable devices.

Other vulnerabilities observable from the CD

1. An attacker can use a host security-scanner, e.g. ISS' System Scanner (TM), to reveal vulnerabilities on the CD. He can then look for other devices on the network with the same vulnerabilities.
2. A security scan of the network, from the CD itself, can reveal other likely targets for the attacker.

LCD-prediction

The first step of the I-ID technique is evidence-collection. It is a systematic means for obtaining the information needed to identify likely-compromised devices. The second step is LCD-prediction. Using the information acquired in evidence-collection, LCDs are identified, along with their degree of likely-compromise.

We believe that assessment of likely-compromise will rely primarily upon subjective human judgement, and not fixed algorithms. The reason being that the input information is incomplete, speculative, and possibly incorrect. Also, the attacker is an intelligent agent who seeks to deceive us and defeat our security measures. As in military battle, there is an art to predicting an attacker's behavior. It depends upon battle experience and wise speculation. However, there are principles and techniques that can be used for LCD-prediction. They are presented later.

There is another consequence from uncertain and incorrect input information. As more is learned about the attack, the collected evidence will be revised, updated and corrected. Our understanding of the attack will be improved, and the LCD-predictions will need to be revised. LCD-prediction is a dynamic and on-going process. The implementation of the I-ID technique needs to accommodate this process.

The identified LCDs will form an LCD-set. Building the LCD-set consists of determining which elements to place in it, determining the degree to which they are likely to be compromised, and then ranking those elements. It is anticipated that the ranking will be categorical (e.g. high, medium, and low) and/or ordinal.

A device can be considered likely-compromised for multiple reasons. The primary reasons are the presence of a known vulnerability, the presence of an asset valued highly by the attacker, or suspicious activity on the device. Other reasons are possible, e.g. information from an informant. The elements of the LCD-set will be uniquely identified by 1) the device, and 2) the reason for likely-compromise. Thus, a single device can have multiple entries in the LCD-set.

The following sections describe some principles and techniques for identifying LCDs:

Economic principles

The results of economic-attribute analysis can be used to identify LCDs:

1. When estimating the attacker's cost, devices with known vulnerabilities were identified. They can be placed in the LCD-set.

2. Some of the devices with known vulnerabilities may have exploitation costs that exceed the attacker's resources. These devices have low, or no, likelihood of attack. For example, a vulnerability's exploitation can require more skill than the attacker possesses.
3. Differences between asset value and cost may indicate degree of likelihood. An asset of little value, but of high cost to obtain, is not likely to be attacked. The converse holds as well.
4. Devices with high asset-value or low cost are likely to be attacked.
5. Basic principles of economics are a guide for prediction: 1) Total costs can not exceed total resources. 2) To engage in an attack, anticipated value must exceed anticipated costs. 3) The attacker will seek to maximize assets and minimize cost.

Opportunistic attacks

Many attackers lack, at the offset, all the network information needed to perform the attack. The "outsider" connecting over the Internet will learn about the network as he attacks it. Even an "insider's" attack may start with incomplete network knowledge.

This lack of information limits attackers in their ability to perform detailed long-term planning. As the attack proceeds, the attacker discovers previously unknown opportunities- both means and ends. For example, the attacker usually can not predict the access (means) he'll gain from cracking a password file. After penetrating a network, the attacker may discover a hitherto unknown file server (ends).

With incomplete network information- as is often the case- the attacker's plan will be dynamic. Both the attack's ends and means are subject to change as more is learned about the network. We refer to such a strategy, or tactic, as ***opportunistic***.

When an opportunistic strategy, or tactic, is being used, economic-attributes provide a useful means for predicting LCD's. For example, the next attack target is likely to be a "weak link" (a device that has low exploitation cost) with attractive value.

Profile attributes

Many of the profile-attributes can be used directly in the prediction of LCDs. Here are some ways they can be applied:

1. Exploits that the attacker has used are likely to be used by him again, if the opportunity arises.
2. The attacker may favor particular operating systems. All other things being equal, devices with those operating systems are more likely to be compromised than devices with other operating systems.
3. If the attacker has been very cautious and stealthful, it's unlikely that he'd implement a noisy exploit, e.g. a packet flood.
4. Traffic logs can be used to determine the likelihood of past attacks. The presence, or absence, of traffic to a device in the LCD-set can indicate the likelihood of a past attack.

IV. IMPLEMENTATION

A database is need for collecting, organizing and working with the information used in evidence-collection and LCD-prediction. The database should accommodate the affects of uncertainty and speculation, i.e. additions, updates, and corrections. For making corrections, it will be helpful to clearly identify uncertain data and data that is deduced from uncertain data. Also, it would be helpful to automate collection of the data, e.g. via a vulnerability-scanning tool.

V. EXTENDING THE I-ID TECHNIQUE FOR ADDITIONAL USES

There are two possible extensions to the I-ID technique. The use of these extensions requires only a small marginal-cost.

One of the extensions was mentioned earlier. The overall objective of the I-ID technique is to identify LCDs. Once identified, the LCDs will need to be investigated for compromise. Investigative techniques are beyond the scope of this paper. Profiling, especially attacker-identification, could be extended for use in investigation.

The second extension provides the sys-admin with risk-assessment for ARNC. Risk is often evaluated as a function of assets, threats and vulnerabilities. The purpose of evaluating threats and vulnerabilities is the prediction of compromise, and the resultant loss of assets. The I-ID technique provides a prediction of compromise for an extended attack. To perform risk-assessment, the only additional information needed is the network's assets, from the perspective of the network owner.

VI. CONCLUSION

The Investigative Intrusion-Detection (I-ID) technique was described. Prior to an attack, an IDS operates in anticipation of a general threat. During an attack, the IDS can deal less in the general and more in the particular- namely, particulars about attackers and attacked devices. A profile of the attacker is developed, using information he reveals about himself during his attacks. Principles from economics are used to predict the attacker's behavior, based on estimates of his asset-appraisal, attack-costs and attack-resources. Likely-compromised devices (LCDs) are identified, using the profile and the economics-based estimates. Knowledge of LCDs is useful for work in attack repair, neutralization and containment.

This is a preliminary and introductory report of work performed for the GIANT project, under DARPA's Information Assurance and Survivability program.

VII. REFERENCES

[FREED] Freedman, David, At Large, 1997, Simon & Schuster
[ICOVE] Icove, David, Computer Crime, 1995, O'Reilly
[OHARA] O'Hara, Charles, Fundamentals of Criminal Investigation, 1994, Thomas

VIII. ACKNOWLEDGEMENTS

Much valuable advice was received from Jim Settle, retired head of the FBI's National Computer Crime Squad. Use of the economics-of-crime was suggested by Dr. E.C. Pasour Jr., NCSU Department of Agricultural Economics. Economics instruction was graciously provided by The Mises Institute. We'd also like to thank Dr. Maurice Godwin, criminal-profiler.

IX. AUTHORS

?? **Jim Yuill** is a Ph.D. Candidate in the Computer Science Department at North Carolina State University. http://www.pobox.com/~jimyuill

?? **S. Felix Wu** is an Assistant Professor of Computer Science at North Carolina State University.

?? **Fegmin Gong** is head of the Advanced Network Research Group at MCNC.

?? **Ming-Yuh Huang** is a Senior Scientist with the Boeing Applied Research and Technology Group.