ELSEVIER

Computer Networks 34 (2000) 671–697

COMPUTER NETWORKS

www.elsevier.com/locate/comnet

Intrusion-detection for incident-response, using a military battlefield-intelligence process

J. Yuill *, F. Wu, J. Settle, F. Gong, R. Forno, M. Huang, J. Asbery

Computer Science Department, North Carolina State University, Box 8206, Raleigh, NC 27695, USA

Abstract

A network device is considered *compromised* when one of its security mechanisms is defeated by an attacker. For many networks, an attacker can compromise many devices before being discovered. However, investigating devices for compromise is costly and time-consuming, making it difficult to investigate all, or even most, of a network's devices. Further, investigation can yield false-negative results. This paper describes an intrusion–detection (ID) technique for incident-response. During an attack, the attacker reveals information about himself and about network vulnerabilities. This information can be used to identify the network's likely compromised devices (LCDs). Knowledge of LCDs is useful when limited resources allow only some of the network's devices to be investigated. During an on-going attack, knowledge of LCDs is also useful for tactical planning. The ID technique is based on the US military's battlefield-intelligence process. Models are constructed of the network, as the battlespace. Also, models are constructed of the attacker's capabilities, intentions, and courses-of-action. The Economics of Crime, a theory which explains criminal behavior, is used to model the attacker's courses-of-action. The models of the network and the attacker are used to identify the devices most likely to be compromised.

Keywords: Computer security; Intrusion detection; Incident response; Military intelligence; Economics of crime

1. Introduction

When a network is under attack, its intrusion-detection system (IDS) faces unique difficulties and opportunities. This paper explores those difficulties and opportunities, and it presents a new intrusion-detection (ID) technique based upon them. The technique is an adaptation of the US military's battlefield-intelligence process, named *Intelligence Preparation of the Battlespace* (IPB) [10]. We have descriptively named the ID technique *Cyber-IPB* (C-IPB).

1.1. The problem

A system-administrator discovers that a hacker has broken into a network device. Unfortunately, for many networks, this discovery is just the tip of the proverbial iceberg. ID tends to be a weak element of network security, giving a hacker opportunity to compromise many devices before finally being detected. Also, network devices often have security *trust-relationships* with other network devices. After compromising one device, the hacker can use trust-relationships to easily compromise additional devices. By the time a successful attack is discovered, many other devices may well be compromised.

* Corresponding author.
E-mail address: jimyuill@pobox.com (J. Yuill).

1389-1286/00/$ - see front matter
PII: S1389-1286(00)00142-0

After discovering one compromised network-device, the system-administrator would like to identify all compromised devices. However, investigating the network for compromise can be a difficult and time-consuming task: (1) devices can be checked manually for telltale signs of compromise, such as strange accounts in /etc/passwd, or suspicious log-file entries, (2) IDSs that run periodically, e.g., Tripwire® [19], can be run immediately, and (3) IDSs can be configured to be more sensitive or to look for specific indications of compromise [1,9]. For networks with more than a few dozen computers, the system-administrators will typically not have time to investigate all, or even most, devices for compromise. In addition, investigating devices for compromise is an uncertain task. The absence of evidence of compromise does not guarantee there is no compromise – investigation is subject to false-negative results.

When network-devices have security trust-relationships, the system-administrator needs to identify and repair *all* compromised devices. If a single compromised device is left on the network, the hacker may be able to continue compromising devices. Also, during an on-going attack, compromised devices must be identified quickly, to minimize attack damage.

The difficulty of identifying compromised devices is exacerbated by the complexity of the network's topology, administration, and use. For the system administrator, the identification of compromised devices can be overwhelming, as the process is resource intensive, urgent, uncertain, and highly complex. In addition, an active threat makes the environment dynamic.

1.2. Current incident-response techniques

In the larger perspective, incident-response (IR) is the overall process for handling the problems of computer misuse, after misuse is discovered. During IR, three measures used to secure a compromised network are: (1) *attack repair*: repairing devices altered by the attacker, (2) *attack neutralization*: fixing vulnerabilities which the attacker has exploited, or which he could exploit, and (3) *attack containment*: temporary measures for limiting an active attack, e.g., blocking all ftp sessions at the firewall. We will refer to *attack repair*, *neutralization*, *and containment* as *ARNC*. [1]

To repair a compromised device, the system-administrator performs, roughly, these tasks: (1) the attacker's active processes are removed, (2) damage from the attack is assessed and repaired, (3) the exploited vulnerability is determined, (4) an appropriate countermeasure for the vulnerability is chosen, based on risk analysis, and (5) the vulnerability is removed by repairing or improving the system. [2]

The identification of compromised devices is an essential part of ARNC, and ARNC is an essential part of IR.

1.3. An overview of the solution

A new ID technique is presented. Its purpose is to assist the system-administrator with the previously described intrusion-detection problems, encountered during incident-response. The objective of the technique is to identify the network devices that are likely to be compromised by the attacker. The devices' degree of likely compromise is also identified. By identifying the devices that are most likely to be compromised, the system-administrator can make effective use of the limited resources for investigating devices for compromise.

As previously mentioned, the ID technique is named C-IPB, and it is an adaptation of a military battlefield-intelligence process. C-IPB provides a systematic method for identifying *likely compromised devices* (LCDs), based on models of the network and the attacker.

A network is attacked by a particular set of individuals. During the attack, each individual reveals information about himself. [3] This information can be used to create models of the attacker' *capabilities* and *intentions*.

[1] *ARNC* is this paper's summary of the measures taken to secure a compromised network. Similar summaries can be found elsewhere, e.g., [20] summarizes the measures as analyze, contain, eliminate, and return [to normal operations].

[2] The repair process is not always this difficult – at times it is possible to just reinstall system software.

[3] This paper's masculine pronouns are used in a gender-neutral manner.

Models of the network are built from the perspective of it being a *battlespace* environment. In the military's battlefield-intelligence process, models are built of the physical, political and sociological aspects of the battlefield [27]. Similarly, in C-IPB, models are built of the network's devices, administration and use. These models reveal the opportunities and difficulties that the network affords for attack. Also, during an attack, the attacker reveals information about network vulnerabilities. This is another source of information for the network model.

Using the models of the network and the attacker, estimates can be made of the attacker's *courses-of-action* (COAs). Two types of COAs are considered: *possible* and *likely*. [4] Estimates of the attacker's possible COAs are based on his capabilities and on the possibilities for attack provided by the network. The possible COAs establish the bounds within which the attacker can operate. Estimates of the attacker's likely COAs are based on his intentions, the network' assets, and the ease with which particular attacks can be carried out on the network.

By identifying the attacker' COAs which lead to compromise, LCDs can be identified, which is the objective of C-IPB. Also, the degree to which a COA is likely indicates the degree to which a device is likely to be compromised.

A real-world IR experience illustrates how an attack reveals network vulnerabilities, as well as the attacker' capabilities and intentions. [5] In 1999, a company' system-administrators discovered that two Linux machines were compromised. Investigation revealed that a single attacker had compromised both machines, using the same exploit on each. He had used the exploit to install his own telnet accounts on the machines. The organization then investigated its other Linux machines to see if they had the same vulnerability and if they were compromised. The attacker' telnet sessions were put under surveillance. The attacker was observed running *IRC bots* – programs for attacking Internet-chat servers [35]. He was not observed attacking other machines in the organization. The system-administrators concluded that the attacker was probably a *script-kiddy*. [6] He appeared to have no particular interest in this company. Once an Internet-chat server detects an attack from a bot, the server blocks all connections from the bot's IP address. Consequently, script-kiddies who run bots need a continual supply of new machines from which to launch their attacks. The system-administrators concluded that the attacker's intention was simply to obtain machines for running bots and that he was not a serious threat to the company.

1.4. Prior work

There is a large volume of publicly published literature on ID research. However, among this literature, there does not appear to be any research which addresses the problem of ID during IR. Unlike ID, little has been published about IR. Within the public IR literature, there appears to be only elementary discussion of the problem of ID during IR. Also, there does not appear to be any publicly published research in the application of intelligence-analysis techniques to computer security. It is likely that the US military has done extensive work in this area, but it is not disclosed. In addition, this paper applies principles from the Economics of Crime [30] to ID. There does not appear to be any publicly published research on this application.

2. The C-IPB process

As just described, C-IPB is an ID technique. Its objective is to identify LCDs. This section contains an overview of C-IPB's four-step process, a description of the military intelligence-process from

[4] In C-IPB, the attributes *possible* and *likely* are typically qualitative, subjective and uncertain. Their use is pragmatic. See Section 2.3, *The nature of the C-IPB process*.

[5] This example is a personal experience of two of the authors.

[6] *Script-kiddy* is a pejorative name for low-skilled teenage hackers who engage in network mischief and vandalism. They attack systems by running *scripts* (i.e., programs) written by skilled hackers.

which C-IPB is derived, and discussion of C-IPB's environmental difficulties of non-determinism and uncertainty.

2.1. C-IPB'S four-step process

In C-IPB, the battlespace is the network environment in which the system-administrator engages the attacker. The attack can be in the present and/or the past. The attacker can be an outsider, or an insider, e.g., an employee. C-IPB is a four-step process:

1. *Define the battlespace environment*. The battlespace's boundaries are defined, an initial evaluation of the battlespace features is made, and sources of intelligence-data are identified.
2. *Describe the battlespace's effects*. An in-depth evaluation is made of the network features which influence attacks, defense, and C-IPB.
3. *Evaluate the threat*. An assessment is made of the attacker's capabilities and intentions.
4. *Determine the threat's COAs and LCDs*. Based on the battlespace and the threat's capabilities and intentions, estimates are made of the threat's possible and likely COAs. The COA estimates are used to identify LCDs – C-IPB's ultimate objective.

Step 4 draws upon principles from the Economics of Crime. The attacker's economic constraints are used for understanding and predicting his behavior. Those constraints are: (1) his *valuation of network assets*, (2) his *costs* for exploitation of vulnerabilities, and (3) his *resources* for attacks.

C-IPB is preceded by a preliminary step, *Establish C-IPB Requirements*. In this step, C-IPB's operational resources and requirements are defined. C-IPB is a subordinate part of IR, which will place tactical and risk-management requirements on C-IPB. Also, the resources for C-IPB need to be specified, as they will typically be less than what are needed for thorough investigation and analysis.

C-IPB's steps will be described in detail, starting in Section 3. The rest of the present section (Section 2) discusses principles from military intelligence which are applicable to C-IPB and also principles which govern C-IPB.

2.2. Military intelligence

2.2.1. The military's IPB process

C-IPB is based on the US military's warfighting intelligence-process, called *Intelligence Preparation of the Battlespace* (IPB). IPB is a process used by both the US Army and Marine Corps (USMC). The process is easily adapted for use in detecting attacks on computer networks.

USMC manuals describe IPB as: "the primary analytical methodology used to produce intelligence in support of the [warfighting] decision-making process" [26]. "It is a systematic, continuous process of analyzing the threat and environment in a specific geographic area to determine and evaluate threat capabilities, vulnerabilities, and probable courses of action. It is designed to support ... planning and decision-making" [27].

The military's IPB process is well developed and battle-tested. The IPB process needed little modification for use in C-IPB, and we sought to preserve IPB's techniques and terminology. To the extent that the C-IPB technique is useful, credit must go to those who developed the IPB process and manuals. C-IPB is based primarily upon the manuals [10,24,26,27].

2.2.2. Intelligence analysis

IPB is based on concepts from intelligence theory [10,24,27]. The concepts are easily adapted for use in C-IPB. A summary of basic intelligence-theory concepts is provided here.

Intelligence is defined as information which provides an accurate and meaningful image of the hostile situation. The lowest level of input in the *intelligence-analysis* process is *data*, which is a collection of facts, e.g., a firewall's log of rejected connections. Data are processed into *information*, which provides meaning to the data, e.g., a summary of the data in the firewall log, consisting of a sorted list of each source address, and the ports it accessed. *Intelligence is created by analyzing and synthesizing data and information to produce knowledge and understanding about the threat and the battlespace*. For example, the firewall-log's summary is analyzed to identify sources which made connections with hostile intent. The ultimate objective of intelligence is to provide information

to the decision-maker, which assists him in making decisions about battling the threat [27].

Military intelligence defines two classes of intelligence. *Descriptive* intelligence describes existing and previously existing conditions. *Estimative* intelligence attempts to anticipate future possibilities and probabilities [24]. The first three steps of the C-IPB process are largely descriptive: (1) *Define the battlespace environment*, (2) *Describe the battlespace's effects*, (3) *Evaluate the threat*. The fourth step is largely estimative, (4) *Determine the threat's courses of action and the likely compromised devices*.

There are two types of descriptive intelligence: *basic* and *current* intelligence. Basic intelligence is general background knowledge about relatively constant conditions. It is gathered in advance of the battle. For example, much of the information collected in C-IPB's Step 2, *Describe the battlespace's effects*, can be collected in preparation for attack. Information about attackers' typical methods of operation is another form of basic intelligence. Current intelligence is concerned with describing the existing condition. The present paper deals primarily with current intelligence.

Fig. 1 is taken from *Marine Corps Warfighting Publication*, *Intelligence Operations* [26]. It depicts the process of converting data into intelligence. The framework of the intelligence process is *analysis*, *synthesis* and *estimation* [27]. Intelligence-analysis theory draws from epistemology, logic, cognition, and perception [14,24,27]. An opportunity for future research is the application of this theory to other areas of ID and IR. For example, in viewing ID and IR from the perspective of this figure, the processes of cognition and judgment have proven difficult to automate. Perhaps significant elements of ID, and IR, rely on processes of cognition and judgment which machines cannot perform. Should this be true, for these elements of ID and IR, automation will be limited to assisting human judgment.

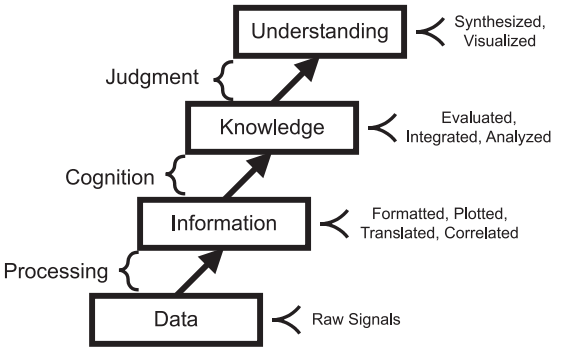

Fig. 1. Converting data into intelligence.

2.3. The nature of the C-IPB process

2.3.1. The need for human judgment

Variations in attackers, networks, and resources-available make battle with an attacker unpredictable. As in military battle, these variations make for an environment which is largely non-deterministic and imprecise. A consequence is that the C-IPB process does not consist of detailed procedure.

C-IPB does, though, provide a systematic and orderly process. It identifies, and applies, useful principles and techniques. Salient characteristics of the battlespace and threat are identified. These principles, techniques, and characteristics are intended to guide the C-IPB analyst, but ultimately he must use his own judgment, insight and cunning in their application.

A similar assertion about the non-deterministic nature of IR was made in one of the subject's earliest papers. In *Responding to computer security incidents*, Dr. Eugene Schultz states, "...it is impossible to specify specific technical procedures for responding to the many types and versions of computer systems within [the Department of Energy]" [33].

A USMC textbook on tactics, echoes the above points: "The tactics involved in warfare are not an exact science. When faced with a tactical problem on the battlefield, you, as a commander, cannot apply a set of rules or a mathematical formula to obtain the ideal solution. You must consider the principles of war and fundamentals of combat that apply to the situation and the factors which affect these general rules. If you fail to recognize and analyze all the influencing factors in an intelligent and orderly manner, you can bring disaster to your own forces" [25].

C-IPB provides an analytical framework for performing ID during IR. The process can be partially automated by databases and graphics-tools. For those experienced in computer-security, parts of the C-IPB technique will be perceived as standard problem-solving. However, as a whole, the C-IPB process is complex, warranting documentation. Also, in the urgency of IR, one needs to react, rather than figure-out what to do.

2.3.2. An environmental constraint: Uncertainty

Information is collected and analyzed, for the purpose of understanding and anticipating the attacker's behavior. 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 human behavior. Also, estimates of past behavior can be uncertain, due to incomplete information. These sources of uncertainty necessitate the use of speculation in analysis and LCD-prediction. In military combat and in entrepreneurial ventures, success depends upon wise speculation about future human behavior. [7] We believe ARNC, including intrusion-detection for ARNC, is similarly speculative.

2.4. Other elements of C-IPB

2.4.1. C-IPB is a continuous process

C-IPB is a four-step process, and the steps are performed roughly in order. However, any step can be re-evaluated or updated as more is learned about the battlespace and the threat. New information can be added, and errors in prior analysis can be corrected. For each step's analysis, we bring to bear all that is known about the battlespace and threat. For example, the third step is *Evaluate the threat*, but some knowledge of the threat is required in the first step, *Define the battlespace environment*.

Feedback and revision in the C-IPB process is a result of: (1) its dynamic environment (e.g., an active threat), (2) the C-IPB responder's continual increase in understanding, (3) the correction of erroneous analysis (e.g., caused by uncertain and deceptive information), and (4) the development of inter-dependent models (i.e., the models of the battlespace and attacker influence each other). [8]

2.4.2. The LCD-set

Collectively, likely compromised-devices (LCDs) form an *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 (e.g., a particular vulnerability).

2.4.3. The dimension of time

C-IPB seeks to identify devices which: (1) have been compromised in the past, (2) are compromised in the present, and (3) will be compromised in the future. So, C-IPB can be applied to a present attack, and, using historical information, to a past attack.

2.4.4. The scope of C-IPB

The scope of 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

[7] Clausewitz goes so far as to state, "in war everything is uncertain". He describes the difficulties of constructing models and theories of war, caused by war's inherent uncertainty and complexity. See [5] Book Two, Chapter 2, "On the Theory of War". The uncertainty inherent in the prediction of future enemy action is described in the USMC doctrinal text on tactics. See [23], Chapter 2, "Military Judgement". Mises states that, "uncertainty of the future is already implied in the very notion of [human] action". He describes the difficulties of constructing models and theories of economics, caused by the uncertainty of future human action. See [28] Chapter VI, "Uncertainty". Chapter XVI, "Prices", describes the entrepreneur's uncertainty in predicting future demand.

[8] Jackson and Cameron's work on software-engineering models is applicable to C-IPB. Cameron states that *decomposition* is often a good process for *describing* an *existing* model. However, *developing* a *new* model is fundamentally a process of *composition* [4].

inside or outside the domain, and it is assumed that there is at least one known, or suspected, attack.

2.5. C-IPB's details

C-IPB's four steps are described in the following sections. C-IPB's preliminary step, *Establish C-IPB Requirements*, is described first. Fig. 2 gives an overview of C-IPB.

C-IPB will be performed by incident-responders. We will refer to them as *C-IPB responders*, or simply *responders*.

3. Establish C-IPB requirements

The major requirements which govern the conduct of C-IPB are: (1) risk management, (2) the resources available, (3) operational policies, and (4) ARNC's ID requirements.

C-IPB is part of, and subordinate to, the process of ARNC. Ultimately, ARNC is subordinate to the overall objectives for the network, e.g., the business objectives which the network supports. To appropriately prioritize his work, and to allocate resources for it, the C-IPB responder needs to understand the higher-level objectives which the network ultimately supports.

ARNC's *risk-management requirements* will direct C-IPB. These include requirements for the protection of network assets. For example, to minimize potential loss, an e-commerce site ranks its top security priority as confidentiality of customer credit-card numbers. Credit-card confidentiality may warrant a relatively large amount of attention from C-IPB, due to the potential loss. Risk-management requirements also include the circumscription of risk in C-IPB activities. For example, after an attacker is detected, rather than stopping him, he can be put under surveillance. The surveillance provides useful information for C-IPB, but at the risk of further damage.

The *resources* for conducting C-IPB are finite, so they will limit the scope of C-IPB. Resources include standard items like C-IPB responders and ID tools. Another resource is assistance from internal and external personnel. For example, during IR, system-administrators can be subordinate to C-IPB responders. Assistance can also be provided by IR-alliances with external networks, and by law-enforcement, e.g., subpoenas for information from other networks.

Legal and organizational *policies* will regulate C-IPB, e.g., counter-attacks are usually illegal.

In battling the attacker, ARNC will have *ID requirements* for C-IPB. ARNC's *tactical requirements* will direct ID. For example, during ARNC the system-administrator is in a race with the attacker, so there will be limitations on the time-available for performing C-IPB tasks. Also, ARNC will have requirements for the location and type of intrusion to be identified, e.g., C-IPB should first discover how the attacker breached the network's firewall. The focus of this paper is intrusion-detection, so ARNC's tactical and risk-management requirements are not elaborated upon further.

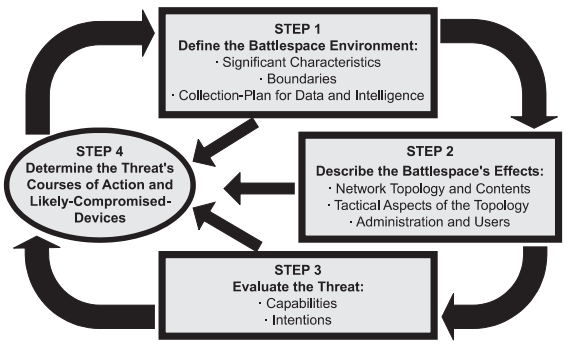

Fig. 2. Cyber Intelligence-Preparation-of-the-Battlespace (C-IPB). It is a continuous process, as shown by the large arrows. Step 4 is derived from Steps 1 to 3, as shown by the small arrows.

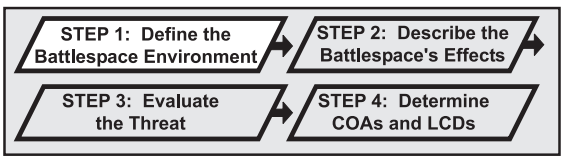

4. C-IPB's Step 1: Define the battlespace environment

The first C-IPB step plans the use of C-IPB, for the present battlespace and attacker. Defining the battlespace environment is a four-part process:

1. Identify significant characteristics of the environment.
2. Identify the battlespace boundaries.
3. Determine the data and intelligence to be collected for C-IPB.
4. Begin collecting the data and intelligence required to conduct the remainder of C-IPB.

4.1. Identify significant characteristics of the environment

The network is not homogenous in its security, architecture, administration, use, nor assets. As a result, it is necessary to identify the primary aspects of the environment which will influence the attacker's COAs and the C-IPB process. Initially, these characteristics are only examined at a high level, to facilitate planning. Further examination takes place in later steps. `Primary aspects of the environment include network topology, organizational realms (e.g., the accounting department), system-administrator realms (e.g., each department has a system-administrator), network uses, and the attackers' capabilities and known activity.`

Details needed for using C-IPB are displayed in a `typewriter font`, as in the preceding paragraph. Readers who are just interested in C-IPB's concepts can skip these details.

4.2. Identify the battlespace boundaries

The *battlespace* consists of the areas-of-operations, areas-of-interest, and areas-of-influence. The *areas-of-operations* (AO) are the realms in which the C-IPB responder has authority and responsibility for ID. The *areas-of-influence* are the realms in which the C-IPB responder has influence over ID. The areas-of-influence are a superset of the AO. For example, when a system-administrator is the C-IPB responder, the AO could be the subnet he administers. The network's other subnets could be his areas-of-influence, by means of cooperation with the network's other system-administrators.

The *areas-of-interest* (AOI) include the AO and those areas beyond the AO from which information and intelligence is required for the C-IPB process. The limits of the AOI are based on the locations from which the attacker can affect, or enter, the AO. Also, if it is anticipated that the AO will be extended into new areas, they can be included in the AOI. The AOI's areas which lie beyond the AO can be larger than the AO. The AOI is not constrained by the ability to acquire information. *Intelligence gaps* will exist in those parts of the AOI for which information is not available.

External networks will typically be part of the AOI, e.g., a trusted external-network connected via a VPN. For publicly accessible network interfaces, e.g., a web server, the AOI is the entire Internet. Clearly, there will be large intelligence gaps. The AOI also includes the sources of known attacks from external networks.

In setting the AOI boundary, the objective is to identify areas with information needed for C-IPB. The information can be actually gathered, or, in the case of intelligence gaps, assumptions will need to be made. Setting the boundaries of the AOI will rely upon good judgment. For example, if a responder is confident that the attacker does not have access to a trusted external network, then that network can be excluded from the AOI.

4.3. Determine the data and intelligence to be collected for C-IPB

There are two parts in the process of planning data-collection and intelligence-analysis. They are performed together and are described below. This planning is for the next three steps of C-IPB, so the implementation of the planning will be clearer after those steps are presented.

4.3.1. Identify the amount of detail required and feasible for C-IPB, with the resources available-including time

The purpose of this step is to plan how the C-IPB efforts should be focused, given the available resources and ARNC's ID requirements. These plans should identify the places within the AO and AI which have the most promising sources of information based on what is known about the battlespace and threat.

The intelligence-analysis process collects data, analyses it and then produces information about

the battlespace and attacker. For the information that is to be produced, the required amount of detail must be determined. The detail will be constrained by the resources available. Also, the network is not homogenous, nor is the attacker's activity within the network. For the different aspects of the battlespace and threat which are analyzed, the degree of detail needed will vary. For example, if an attacker is known to only compromise Linux devices, then more detailed information may be needed about subnets with Linux devices than subnets without them. Aggregation is one way to vary detail. A subnet without Linux devices might be adequately described as having, "about 100 Solaris hosts". However, for a subnet with Linux devices, details would be needed about each Linux device's version, servers, and exploitable trust relationships (e.g., file sharing).

For data which does exist, and is useful, the C-IPB responder will need to identify that which is feasible to analyze. For example, there may be distributed log-records which contain useful information, but the cost of collecting and analyzing them may be prohibitive.

4.3.2. Evaluate existing sources of useful data and identify intelligence gaps

The existing ID literature identifies sources of information which are useful for ID, e.g., [1]. The primary sources are device log-records. Routers, firewalls, and hosts keep such records. IDSs are another source of ID data. So-called out-of-band sources are also useful. For example, users and system-administrators can detect suspicious network activity and report it.

The C-IPB responder should identify the useful data which: (1) does exist, (2) might exist, and (3) does not exist. For data which might exist, searching for it incurs a cost, e.g., workstations can be examined to determine what logging is performed. For data which can't be obtained, wise assumptions must be made.

For data which does not exist, it may be possible to start obtaining it. New ID mechanisms can be installed. Existing ID mechanisms can be configured to collect more data, making detection more likely, but often at the cost of a greater signal-to-noise ratio. An IDS can be configured to be more sensitive in its detection, but usually at the cost of increased false-alarms. These added costs may be prohibitive during normal operations, but justifiable during incident-response. Also, during C-IPB, knowledge of the attacker can be used in deploying and configuring ID mechanisms, increasing the likelihood of detection.

Throughout the C-IPB process, some of the intelligence gathering and analysis can be done before an attack occurs. Such preparation increases the C-IPB responder's speed and his depth of his understanding. A lesson from guerilla warfare is that superior knowledge of the local terrain can provide a decisive advantage to an indigenous force [18].

4.4. Begin collecting the data and intelligence required to conduct the remainder of C-IPB

The C-IPB responder can begin collecting and analyzing the data that was identified earlier in this step. The C-IPB process is *dynamic* – each part of the process is subject to revision as more is learned about the battlespace and the attacker. The confirmation or refutation of assumptions is a significant cause of revision. Confirmation provides greater certainty. It allows the C-IPB process to be more focused, as fewer alternatives and contingencies need to be considered. Reasoning and actions based on refuted assumptions will need to be corrected. There will almost always be needed information which is unknown or uncertain, so assumptions are a necessary part of the C-IPB process.

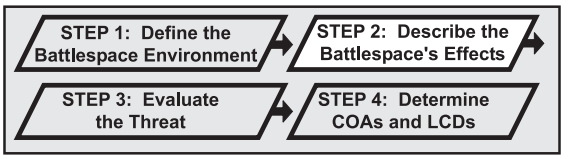

5. C-IPB's Step 2: Describe the battlespace's effects

5.1. Introduction

The purpose of this step is to determine how the battlespace environment affects the capabilities

and COAs of both the attacker and of C-IPB. The primary aspects of the battlespace are:

- where, in the battlespace, the attacker can be;
- tactics, techniques and attacks that the attacker can use in the battlespace;
- opportunities for detecting the attacker.

Another important aspect of the battlespace is the strengths and weaknesses which it affords both the attacker and C-IPB. Understanding these strengths and weaknesses enables the C-IPB responder to pit his strengths against the attacker's weaknesses. This understanding also reduces the attacker's opportunity for surprise.

The battlespace's effects are determined by analyzing these broad factors of the battlespace:

Network topology
Tactical aspects of the topology
Compromised devices and known vulnerabilities
System administration
Network users

These factors are described in the following sections. The factors' specific features are often relevant, but not always. Also, the description is not exhaustive. As mentioned earlier, C-IPB does not lend itself well to detailed cookbook-like processes.

The C-IPB responder has limited resources. He will need to be creative and wise in his use of them. For example, at times it will be sufficient to generalize a battlespace factor, rather than describe it in detail, e.g., "the subnet has about 100 Linux workstations".

In analyzing the battlespace, the C-IPB responder should make use of all that is currently known of the battlespace and the threat, especially his capabilities, intentions and likely COAs. Initially, little may be known about the threat. In this case it is only possible to identify aspects of the battlespace which affect many of the possible threats and COAs, or which affect the most likely threats and COAs. When little is known, the analysis should be broad, providing orientation needed for focusing further analysis. This step of the C-IPB process can be revisited and further analysis performed when more is known about the threat and likely COAs.

Since the battlespace is not homogenous it is necessary to identify and focus on those areas which are most important – the areas where the attacker is most likely to be found, and the areas most relevant to C-IPB's risk-management and ARNC requirements. Typically, these will be the areas which are accessible from, or which have access to, the compromised devices. Also, areas that are not likely to be compromised, or used by the attacker, can be excluded from consideration, effectively reducing the battlespace. The larger the network, the more pressing it is to reduce the battlespace – lest the amount of information become overwhelming.

In military battlefield-intelligence, graphics are the preferred means for conveying an image of the battlespace [26]. A map of the battlespace is created for each of the battlespace characteristics, e.g., obstacles and observation points. The maps are then overlaid to visualize the combined effects of the characteristics. For C-IPB's battlespace factors, a similar approach could be used. The graphical representation of C-IPB's models is left as a topic for future research. However, useful graphics can still be created using ad hoc techniques.

The following sections describe the battlespace factors which were just introduced.

5.2. Network topology

This section summarizes the elements of the network topology which affect the battlespace. (They are well known to computer-security practitioners and in the interest of space are not elaborated upon.)

A map of the network architecture and components provides a framework for analysis of the battlespace's effects. The primary features include networking devices and their links (e.g., routers, switches and hubs), network-management facilities (e.g., SNMP managers and agents), host systems (e.g., hardware, operating system and network servers), and host content (e.g., application software and data).

The network's security measures provide the basis for protection and ID. Primary features

include network-security measures (e.g., router filters, firewalls, VLANs, and encrypted channels), device-security measures (e.g., passwords for access control), ID systems and resources (e.g., logs on routers, firewalls and hosts, and IDSs for networks and hosts), and intrusion-response systems (e.g., automated updates of firewall-rules for detected attacks).

The network's content is another important feature of the topology, e.g., application software and data. It has value to both the network owner and the attacker.

5.3. Tactical aspects of the topology

The *tactical aspects of the topology* [9] are those aspects of the topology which govern attack and defense. By knowing the attacker's tactical options within the topology, we can know the opportunities he is afforded; we can predict where he is likely to be found and how he is likely to be working. On the other hand, by knowing C-IPB's tactical options within the topology, we can effectively deploy detection resources.

The tactical aspects of the topology are described below. They are: (1) Observation and opportunities for stealth, (2) Zones-of-attack and cover, (3) Network-path obstacles, (4) Avenues of approach, (5) The attacker's capabilities for collecting intelligence, and (6) Key network tactical-assets.

Two characteristics of the attacker's tactical options are: *exploitability* and *sustainability*. Exploitability is the degree of difficulty in appropriating the use of some network feature. Sustainability is how long an exploited vulnerability can be compromised. Sustainability is influenced by the likelihood of detection, and also by the system-administrator's opportunities for repair and neutralization of the exploit. For example, a readily detected web-defacement is sustainable if the system-administrator can be prevented from repairing the defacement and from neutralizing the exploited vulnerability.

Some elements of tactical-analysis depend upon knowledge of the attacker's capabilities, intentions, and likely COAs. For example, once the attacker's likely targets are known, then the possible routes to the targets can be identified.

5.3.1. Observation and opportunities for stealth

Observation is the ability to see information on the network or on a device. There are observation points which are of tactical use, for the attacker or for C-IPB. For example, Ethernet hubs can be sniffed. Other techniques for observation are port scans (e.g., using the tool nmap [13]) and vulnerability scans (e.g., using the tool SATAN [8]). An example of a tool used for observation on a device is BO2K [7]. It is installed via a Trojan Horse, and it permits a remote user to monitor and control a computer running MS Windows®.

Visibility at an observation point should also be considered. Data visibility can be impaired by various means, e.g., encryption, foreign languages, or application-specific data-formats such as Lotus Notes'®.

A *detection point* is a point along a path that data must pass-through to reach the destination. (Conversely, it is a point which cannot be bypassed in reaching the destination.) Detection points can be useful for observation, especially along avenues-of-approach (described later). Detection points can exist as part of the network's normal configuration, or they can be created for surveillance of the attacker.

Stealth is avoidance of detection. The network can provide opportunities for stealth, which are of tactical use for the attacker or for C-IPB. An attacker can be detected by an IDS or by otherwise doing something conspicuous. Gaps in a network's IDS provide opportunities for stealth. Areas where an attacker can operate stealthily are likely locations for attacks which have been sustained and undetected.

5.3.2. Zones-of-attack and cover

Zones-of-attack are those areas that can be attacked from a given position. For example, behind

[9] This concept is an adaptation of IPB's *Military Aspects of the Terrain* [27].

firewall F–V lie two subnets – subnet S (Secure) and subnet V (Vulnerable). Subnet S is protected by its own firewall, F–S. An attacker has compromised F–V, but not F–S. Subnet V is his zone of attack, through F–V.

Cover provides an attacker protection from ARNC. For example, an attacker can block a system-administrator's remote access to a device, thereby acquiring cover until the system-administrator has physical access to the device. Cover is useful when an attack is to be sustained after detection. For example, cover is needed to sustain a conspicuous web-page defacement.

5.3.3. Network-path obstacles

An *obstacle* in a network-path is any network feature which stops, impedes or diverts the attacker in his attempt to get from one point to another in the network. The evaluation of obstacles helps to identify mobility-corridors (described later). An obstacle can be a device which forwards data on the path, e.g., firewalls and network-address-translation (NAT) devices. An obstacle can be a device which dissuades the attacker from using a network path, e.g., an ID device.

Obstacles can be created intentionally, e.g., a firewall. They may also be unintentional, e.g., a NAT device. An obstacle's effect on network passage can be divided into three categories:

- *Unrestricted*: Network paths which are free of any restrictions to data flow, e.g., a network attached directly to the Internet.
- *Restricted*: Network paths whose data flow is hindered to some degree. For example, a firewall blocks access from the accounting department's subnet, but not the engineering department's subnet.
- *Severely restricted*: Network paths whose data flow is impossible or impractical.

The effects of multiple obstacles can be assessed collectively, along a path. If a path contains many obstacles with restricted passage, the overall effect could be a severely restricted passage.

An obstacle's effect on passage can vary, depending upon such things as:

- the direction of traffic, e.g., a firewall may let traffic out, but not in,
- the specific path taken, e.g., in the example above for restricted passage, if the attacker is already in the engineering department, the passage is unrestricted,
- the needed bandwidth, e.g., a low-bandwidth link may be severely restricted for a packet-flood attack, but unrestricted for a telnet session,
- the type of traffic, e.g., a proxy firewall may only permit connections to web servers,
- the attacker's skill, e.g., what is impossible for a low-skilled attacker may be easy for a skilled attacker.

A cautionary note is warranted here – *in both war and computer security, many successful attacks have been carried out by a clever attacker who traversed terrain the defender assessed impassible*.

5.3.4. Avenues of approach

Avenues of approach (AAs) are routes the attacker can take to reach his objectives. Determining AAs requires some understanding of the attacker's likely COAs. In particular, it requires some understanding of where he is coming from and where he is going to. Determining the threat's COAs is the fourth step in the C-IPB process. The attacker's tactics and techniques will influence his choice of AAs. Assessment of these tactics and techniques is made in the third step of the C-IPB process, *Evaluate the threat*. The present step can be revisited when more is known about the threat's COAs and tactics.

There are two steps for developing AAs: (1) identify mobility corridors, and (2) link mobility corridors to form AAs. *Mobility corridors* are the paths that the attacker can potentially traverse in the network. They are primarily the paths established by network-path obstacles and by routing devices.

It can be helpful to distinguish between two means of entry to a network, or a particular part of a network. *Intended paths-of-entry* are the means designed for entry, e.g., an Internet connection to a firewall's DMZ. *Unintended paths-of-entry* are those which are explicitly not intended, or which just happen to exist. An example of the former is a modem installed on a workstation, in violation of an organization's security policy. An

example of the latter is a dual-homed host which unintentionally provides a path into a subnet.

Mobility corridors can be classified according to the degree of obstruction encountered along the path – unrestricted, restricted, or severely restricted. If enough is known about COAs and about the mobility corridors, they can be ranked in order of likely use. Attackers will favor paths with fewer obstacles and with less likelihood of detection.

Mobility corridors can be linked together to form AAs. Features of the attacker's most appealing AAs are: (1) available paths-of-entry, (2) directness to objective, (3) few obstacles, (4) low likelihood of detection, and (5) sustainable access. As with mobility corridors, AAs can be classified according to degree of obstruction and ranked in order of likely use.

5.3.5. The attacker's capabilities for collecting intelligence

Some sources of information about this particular network can aid the attacker. Information intended for public, or internal, dissemination can reveal information about the network. Examples of such internal information are documentation of the network topology and user-manuals for internal systems. Books on hacking techniques describe other means an attacker can use for collecting intelligence [2].

5.3.6. Key network tactical-assets

Key network tactical-assets are devices or paths whose control or access affords a marked tactical advantage. For the purpose of ID, it is useful to identify the attacker's key network tactical-assets, as he is likely to be found there. A path that is the sole path to an objective can be a key network tactical-asset. Other key network tactical-assets are detection points and observation points.

5.4. Compromised devices and known vulnerabilities

5.4.1. Compromised devices

During ARNC, there are two types of compromised devices (CDs): Known CDs and Likely CDs. Compromise can be of any security attribute, e.g., confidentiality, integrity, availability, authentication, access control, or non-repudiation. CD's need to be identified, along with their tactical use for the attacker, e.g., entry point or observation point.

Known CDs (KCDs) are devices which are known to be compromised. A device which is the source of an attack also will be considered a KCD. Note that this includes devices which are outside the AO and devices to which the attacker has legitimate access.

Likely CDs (LCDs) are devices which are likely to be compromised, either in the past or future. The fourth step of the C-IPB process is LCD-prediction, based on the attacker's capabilities and intentions. Other reasons for considering a device to be an LCD are observation of suspicious activity on the network and reports from external sources about attacks originating from the AO.

KCDs can be repaired and thus not compromised. LCDs can be reclassified as not likely compromised. Lists of KCDs and LCDs are dynamic and need to be updated if a device is no longer a KCD or LCD. It can also be useful to keep a record of the devices which were previously KCDs and LCDs. For example, if a device is compromised a second time, such records would reveal that repair of the first compromise was inadequate or that the device is highly valued by the attacker.

5.4.2. Known vulnerabilities

Vulnerabilities in network-attached devices can be classified as intentional and unintentional. Trust relationships are a primary source of *intentional* vulnerabilities, e.g., file sharing. *Unintentional* vulnerabilities can be in preventive security-measures or in IDSs. The former are usually the result of system-software bugs or configuration errors. A network or host vulnerability-scanner can be used to find them. For an example of an IDS vulnerability – some network IDSs are unable to observe all network traffic during heavy traffic loads.

Vulnerabilities can be further classified by the type of security compromised (confidentiality, access control, etc.), the potential loss from compromise, and the difficulty of compromising the vulnerability.

Depending upon the network, there may be many more vulnerabilities than there are C-IPB resources available to investigate them. The scope of the investigation can be narrowed. The more relevant vulnerabilities are those which, (1) are accessible from KCDs or LCDs, (2) fall within the scope of the attacker's capabilities and intentions, and (3) are involved in his likely courses-of-action. The latter two items are the subject of C-IPB's third and fourth steps.

Vulnerable devices can be detected by using a vulnerability scanner, e.g., SATAN. Another way to find vulnerable devices is to use the sets of devices that are known to be compromised (KCDs) and that are suspected to be compromised (LCDs). We will refer to the combined sets of KCDs and LCDs as the *compromised device* (CD) set, and we will refer to elements of the CD-set as *CD*s. Devices in the CD-set can be put under surveillance, e.g., by sniffing telnet sessions. The attacker' traffic to and from a CD may reveal vulnerable devices.

There is a third way to find vulnerable devices. *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.* Attack-techniques are documented in many sources [2,16,22]. The following sections overview the primary attack-techniques which the C-IPB responder can use to identify vulnerable devices.

5.4.2.1. System software and configuration errors. A common means of compromise are 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. When the system-administrator cannot identify the exploited vulnerability, it may be possible to identify the system-software or configuration that was likely exploited, e.g., the CD could be running a notoriously insecure server. It would be reasonable, then, to be suspicious of other devices with the same such system-software or configuration.

5.4.2.2. 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. In computer-security jargon, this configuration is humorously called, *crunchy on the outside*, *chewy on the inside*. [10]

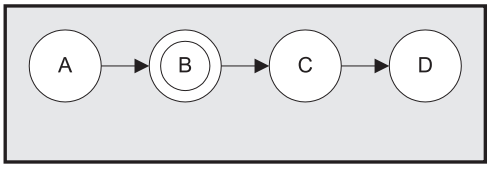

Fig. 3. Graph of exploitable trust relationships.

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, rhosts permissions, and anonymous ftp.

Vulnerable trust relationships can be used to identify: (1) devices which may have been attacked from the CD, and (2) devices from which the attacker accessed the CD.

A directed graph can be used to map exploitable trust-relationships. Nodes in the graph represent devices. Arcs represent exploitable trust relationships. An arc pointing from node A to node B represents "B trusts A, and B can be compromised from A". Paths of trust relationships can be identified. In Fig. 3, node B is a compromised-device. A could have been the means by which B was compromised. C may be compromised now or in the future. If C is compromised, D can be compromised. [11]

5.4.2.3. Exploitable information. A compromised-device may contain information whose disclosure renders other devices vulnerable. For example, a sniffable network connection can provide telnet user-ids and passwords. Mail folders or text files can contain sensitive information such as passwords or maps of the network topology.

5.4.2.4. Other vulnerabilities observable from the CD. An attacker can use a freeware, or commer-

[10] Ref. [32] attributes this humorous description to Bell Lab's Bill Cheswick.

[11] Ref. [6] describes the use of directed graphs for assessing protection rights.

cial, host security-scanner to reveal vulnerabilities on the CD. He can then look for other devices on the network with the same vulnerabilities. A network security-scanner, e.g., SATAN, can be run from the CD itself, revealing other likely targets for the attacker.

5.4.3. System administration

System-administration is the implementation and operation of the network. In practice, it has a controlling influence over security and over C-IPB itself. There are realms of administrative control. For example, individual departments may have their own system-administrators, and the corporate IT department its own system-administrators. For the realms of administrative control, identify those aspects of network administration which affect the attacker's operations and which affect C-IPB operations. Primary aspects are:

- resources and abilities for secure administration, e.g., the system-administrator wants to maintain a secure network, but he has little time for, and training in, security,
- prior performance of security efforts, e.g., the system-administrator's last security problem was over two years ago,
- resources available for assisting with C-IPB, e.g., the system-administrator is cooperative, but he can only spend 1 hour per day assisting the C-IPB responder,
- security policy and its actual implementation, e.g., the system-administrators have a good security policy and implement it effectively.

5.4.4. Network users

In practice, the network users play a key role in security. Analysis of users can be made collectively, e.g., for a department, or on an individual basis. The analysis should identify those aspects of user behavior which affect the attacker's operations and which affect C-IPB operations. The primary aspects include:

- security policies and actual practice, e.g., there is a good policy on paper, but users are not diligent about following it,
- security knowledge and training, e.g., twice a year, users are given instruction in the need for, and practice of, security,
- disposition and attitude toward security, e.g., employee morale-problems indicate that insider attacks are likely.

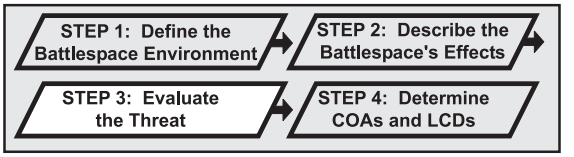

6. C-IPB's Step 3: Evaluate the threat

6.1. Introduction

In this step, we seek to develop a model of the attacker. In particular, we seek to learn his *capabilities*, *intentions*, and *personality* traits, which govern his behavior on the network. This analysis is derived from information about *what the attacker has done*. Knowledge of the threat's capabilities, intentions, and personality traits provide the basis for developing a model of his *disposition* and for discovering his *vulnerabilities to detection*. First, we'll consider analysis of a single attacker and then *multiple attackers*.

An ever-important tactic is pitting strength against weakness. During analysis of the attacker, it is important to identify the attacker's strengths and weaknesses for being detected. Also important are C-IPB's strengths and weaknesses for detecting the attacker.

These components of threat-evaluation are presented in the following sections:

What the attacker has done,
Capabilities,
Personal traits,
Intentions,
Multiple attackers.

As mentioned earlier, details needed for using C-IPB are displayed in a typewriter font. Readers who are just interested in C-IPB's concepts can skip these details.

6.2. What the attacker has done

Information about the threat and his activities is gathered to build the previously mentioned models. The primary sources of information are KCDs, LCDs, and IDSs. Techniques for obtaining attack-evidence from these sources are documented in the ID literature [1,9].

A little used but potentially powerful means of obtaining information about an attacker are so-called honey-pots. They are network devices which are designed to lure would-be attackers. A honey-pot's purpose is: (1) keeping the attacker away from other devices, (2) intrusion detection, and (3) surveillance [1].

For C-IPB, not all details of attacker activity need be recorded, just those which are useful for the end-goal of identifying LCDs. The recording of attacker activity also needs to be cost-effective and within the resources available.

Attack evidence varies from being certain to being highly speculative. The evidence can be incorrect or incomplete. Analysis based on uncertain, incomplete, or incorrect information will need to be revised when better information is acquired. Procedures and techniques for accommodating such revisions are needed throughout the C-IPB process, and it is left as a topic for future research.

Information about the attacker and his activity can be divided into four categories: (1) *Attacks*, (2) *Knowledge of the network*, (3) *Use of devices and data*, and (4) *A priori knowledge*. They are described in the following sections. For each of the four categories, the following *general attributes* are often worth recording:

- the time of the activity,
- patterns of behavior, e.g., based on time, tactics, or network access,
- for network activity, the source and destination addresses, and the path taken,
- devices accessed - hardware, operating system, servers and applications,
- data accessed,
- tools and techniques used,
- files left on systems, [12]
- information that can be used for intrusion-detection, either for investigating LCDs for compromise, or for use by IDSs. Examples are attack signatures, and attack indications and warnings.

6.2.1. Attacks

The *general attributes*, listed above, can be used to describe the attack. In addition, taxonomies of security intrusions and faults can provide useful generalizations. [3,21] Other relevant attributes of an attack are:

- degree of success,
- type of security compromised: confidentiality, integrity, availability, non-repudiation,
- vulnerability compromised,
- exploit used.

6.2.2. Knowledge of the network

The attacker's activity includes acquiring information useful for attacks. This information constrains and guides his behavior. The attacker's sources of information about the network are discussed in Section 5.3.5. *The attacker's capabilities for collecting intelligence*.

6.2.3. Use of devices and data

Once an attacker gains unauthorized access to a network device, what he does with it can reveal his capabilities and intentions. His activity can

[12] Ref. [36] notes that, "Intruders often leave all sorts of files on the systems that they compromise". Generically, they are named *remnant files*. Remnant files which may be malicious code are named *artifacts*.

be observed by placing the compromised device under surveillance. This tactic is discussed in the IR literature [20]. Logs and artifacts can also reveal his activity. Examples of activity which reveals capabilities and intentions are: (1) sending messages, e.g., via Internet-chat or e-mail, (2) storing data, or running servers, on the device, (3) connection hopping for account laundering. In addition, the attacker's authorized access to network devices can reveal his capabilities and intentions, e.g., his reconnaissance activity on a public web-server.

An attacker's use of compromised data can reveal his capabilities and intentions. For example, a company's competitor consistently wins competitive bids by a small margin, so it is suspected that the attacker has compromised the confidentiality of the bidding data.

6.2.4. A priori knowledge

Knowledge of the attacker may be obtained from sources other than his present attacks. For example, when the attacker is a known insider, the organization will have information about him. Incident-response teams or law enforcement may have information about a known outsider, from his attacks on other networks.

6.3. Capabilities

The things the attacker does reveal his capabilities. The specific attributes considered are his:

- *Abilities*: what the attacker can do,
- *Method of operation*: what the attacker does,
- *Knowledge*: what the attacker knows about the network,
- *Possessions*: what the attacker has compromised,
- *Exploitable vulnerabilities*: what the attacker can compromise.

A useful source of information on attacker capabilities is Steve Romig's document entitled, *State of the Hack* [31]. He profiles the script-kiddies that he has battled at Ohio State University over the course of a year. The document provides background information on the script-kiddy's demographics, personality, skill, training and techniques. From the perspective of intelligence analysis [31], provides *basic intelligence* on attackers. There are a number of books about hackers, and hacker techniques, which can also be used for basic intelligence [2,11,12,16,22,34].

In intelligence analysis, when needed information is not available, assumptions must be made. When *current intelligence* is insufficient, *basic intelligence* can provide a useful basis for assumptions. For example, if an attacker appears to be a script-kiddy, it may be reasonable to assume he will act like a typical script-kiddy.

In addition to understanding what the attacker's capabilities are, it is also important to understand what he perceives his capabilities to be. The attacker will attempt what he thinks is possible, not necessarily what is possible. The attacker's over-estimation of his ability can result in fruitless effort or in detection. Under-estimation can result in missed opportunity or unnecessary delay.

6.3.1. Abilities: what the attacker can do

Attributes of the attacker's abilities are:

(a) The attacker's *computer-skill* is a function of his general technical abilities. Two areas of skill are the operating systems he knows and his programming ability. For example, the attacker may demonstrate strong Unix skills and weak Windows skills. He may be able 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 routine manner.

(c) The attacker's *tenacity* describes his persistence. A good example of tenacity is the hacker Matt Singer, as described in the book *At Large* [12]. 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. Ref. [12] describes Matt Singer as unorganized. Some attacks require meticulous record-keeping, and they can only be accomplished by a highly organized attacker.

Computer-skill and attack-skill can be augmented by the attacker's *advisors*. For example, [12] states that 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.

Romig describes the training process of script-kiddies. He has observed that they work in groups and actively teach each other. They are typically "intelligent", but not "computer geniuses". The more talented members provide innovation and leadership. Groups appear to last a year or two. There are generations of groups, and some of a group's techniques can be traced back to its ancestor groups. Knowledge of an attacker's technique may be used to identify the groups to which he belongs. Knowledge of his group's techniques can be used to estimate other techniques the attacker may use [31].

Analysis of ability can include estimates of what the attacker cannot do. For example, it is useful to know both the attacker's strengths and weaknesses for avoiding detection.

6.3.2. Method of operation: what the attacker does

The attacker's *method of operation*, or *MO*, describes what he does, in general terms. In the field of criminal investigation, the MO 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 [29].

The scope of consideration can include both the attacker's tactical and strategic behavior. Tactics are the techniques used during an individual attack. Strategy is the use of multiple attacks to achieve an overall objective. [13]

[13] This follows, roughly, Clausewitz's definitions of tactics and strategy [5].

Some of the MO attributes 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. This is a function of his concerns over being detected, and his perception of the likelihood of being detected.
- Attack technique, e.g., a port-scan followed immediately by an exploit.
- Time spent on the network, both duration and patterns-of-occurrence. This will help predict the frequency at which he works, the volume of his work, and when he is present. Knowing the proportion of time he spends on reconnaissance, attack and system — use is helpful for predicting behavior and motive. It may be possible to determine how much time he has available for attacking, e.g., only evenings and weekends.
- Use of a device or data, once access is obtained.

Patterns in the use of techniques and in the times-of-occurrence can be useful for prediction. These patterns can reveal the attacker's *preferred tactics*. As mentioned earlier, attacker's who are members of a group may use similar techniques and thus have similar MOs. There are several popular books and Internet-distributed documents on hacking technique [2,16,22]. If an attacker's technique is taken from these sources, the sources can provide a means of knowing his MO.

The attacker's MO may depend upon his motive for attacking a particular device. Romig has identified three attacks in which there is a correlation between attack technique and the motive for choosing a victim device. For a *test attack*, the attacker has obtained a new exploit and indiscriminately chooses a vulnerable victim on which to test it. For a *directed break-in* the attacker has a specific goal. He will bring a collection of tools, unlike the testing attack. Also, he is more likely to

act stealthily. For a *convenience break-in*, the attacker is provided unauthorized access by another hacker who has already compromised the device. Access is traded among hackers and even publicly broadcast [31].

6.3.3. Knowledge: what the attacker knows about the network

The attacker's *knowledge of the network* will limit and influence his decision-making. The second step of the C-IPB process, *Describe the battlespace's effects*, analyzes the elements of the battlespace which influence the attacker and C-IPB. The attacker's knowledge of the network can be described in terms of what he knows about the battlespace's effects. Knowledge of the attacker's *misperceptions* can be useful. 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.

In addition to knowledge of the network, the attacker's *understanding* should 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.

6.3.4. Possessions: what the attacker has compromised

The attacker's compromised devices can provide opportunity for compromising yet other devices. Identification of present and past KCDs and LCDs is needed for determining the other devices he is capable of compromising. Attributes of compromised devices were discussed in Section 5.4 *Compromised devices and known vulnerabilities*.

6.3.5. Exploitable vulnerabilities: what the attacker can compromise

In assessing the attacker's capabilities, we would like to identify the devices which could be compromised, but are not presently known to be compromised. Identification of exploitable vulnerabilities was discussed in Section 5.4.2 *Known vulnerabilities.*

Combining knowledge of the network-devices' vulnerabilities and the attacker's capabilities, an estimate of his *exploitation-costs* can be made. Exploitation-costs are for a particular device and vulnerability. The costs are relative to the attacker's capabilities. Costs include: (1) skills needed, (2) time required, (3) difficulty, and (4) likelihood of detection. Without an exploitable vulnerability, a device's exploitation-costs can be considered infinite.

6.4. Personality traits

The attacker's personality traits may be revealed during attack, or known a priori, as with an insider. We are interested in traits which govern his behavior on the network and which can be used for C-IPB:

- *Judgment* summarizes the degree to which the attacker thinks clearly. Judgment can be impaired by vices like greed, arrogance, obsession, and vengeance. As described by [12], Matt Singer is an extreme example of obsession.
- *Age* and *maturity* influence goals and judgment.
- *Morality* governs the degree to which he is willing to inflict loss. Romig observed that the attackers on his network were often involved in other crimes [31].
- *Patience* is needed for stealth and the pursuit of long term goals. Impetuous or reckless behavior is vulnerable to detection.
- *Cautiousness* influences the risk he is willing to take and the precautions he takes.
- The attacker's *culture* may influence his attack behavior. For example, attackers from Tokyo and New York City might differ in their willingness to inflict loss. *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 [29]. Psycholinguistics and cultural influences on attackers are areas for future research.

6.5. Intentions

The things the attacker does reveal his intentions with the network. The specific attributes considered are his:

- *Motives*: what the attacker wants,
- *Appraisement*: how the attacker values network assets.

In general, it is easier to profile behavior than it is to profile psychological attributes like knowledge, personality, motive and asset-appraisement. Behavior can be objectively observed. Psychological attributes tend to be known through subjective speculation. [14] Assessments of psychological attributes can be strengthened by using multiple sources of analysis for corroboration, e.g., skill, MO and knowledge.

Patterns of behavior can indicate intention, e.g., repeated port scans of a particular device indicate the attacker has a special interest in it. Also, an unfolding plan can reveal motive.

6.5.1. Motives: what the attacker wants

Using information about the attacker's activity, we seek to determine his motives and his goals. Two elements of motive are: (1) what the attacker wishes to *acquire*, and (2) what the attacker does not wish to *lose*.

Motive is the primary distinction used in the FBI's categorization of types of attackers: crackers (benign network explorers), vandals, and thieves [15]. Crackers are motivated by the process of hacking itself. Vandals and thieves are more motivated by a particular objective. Landreth offers another categorization: Novice (kids with short attention spans), Student (college-age students curious about security), Tourist (just looking for something interesting), Crasher (delights in denial of service), Thief (serious, knowledgeable, and criminal). Landreth's categorization is partly based on strategic motive [15]. Romig's experiences indicate that the primary motive of script-kiddies is the attainment and preservation of status among peers [31].

Motive is revealed by the attacker's choice of targets. A thief seeking confidential data would have no interest in the contents of a publicly available web-server. However, those contents would be of great interest to a vandal or script-kiddy.

Motive influences the type of attacks that are performed. Romig observed four types of incidents: (1) denial of service, (2) convenience, (3) testing, and (4) directed attack [31]. They were described earlier in Section 6.3.2 *Method of operation: what the attacker does*. Each attack reflects a different, and distinct, motive.

Goals are the specific objects sought in the fulfillment of motive. There can be intermediate goals and end goals. The intermediate goals are tactical, providing the means for achieving end goals. The purpose for compromising a firewall may be to gain access to a file-server behind it. Intermediate goals are related to the attacker's plans, which are considered in the next C-IPB step, *Determine the threat's courses-of-action and the likely compromised-devices*.

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, he is not likely to expend a lot of resources obtaining an asset that is available elsewhere at lower cost. Included in his expenses is stealth. If the attacker does not act stealthily, he may not be interested in this network in particular.

The attacker is also motivated by that which he 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.

[14] The difficulty of assessing psychological attributes is discussed in the following sources. Three different fields are cited: (1) *military-theory*: [24], Chapter 2, "Capabilities versus intentions", (2) *criminal investigation*: [17], Section 2.5, "The Personality to Behavior Confusion", (3) *economics*: [28], Chapter XVI, Section 2, "Valuation and Appraisement".

6.5.2. Appraisement: how the attacker values network assets

For devices on the network, we can *appraise* their 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., Icove's crackers, criminals and vandals). During an attack, analysis of the attacker's motive can be used to create a more accurate appraisement.

For example, consider a network that contains a technically fascinating client-server system and a simple web server. For a high-skilled cracker (Icove's "benign-network-explorer"), 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.

In most cases it would not be possible, or necessary, to appraise every network device individually. Devices can be appraised by type, e.g., routers, workstations, servers. Appraisement is only needed in those parts of the network which the attacker has accessed or will access. Of greatest importance to C-IPB are *high-value-devices* (HVDs). It is useful to ordinally, or categorically (e.g., high, medium, low) rank appraised devices.

6.6. Multiple attackers

Reports from major hacking cases reveal that hackers often work in small and loosely formed groups [12,31,34]. 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, e.g., by pooling disparate skills.

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, analysis of the group's capabilities, intentions, and courses-of-action can be a means for intrusion detection. Ultimately, groups are made of individuals. Analysis of individual attackers will be the basis for analysis of a group.

Loosely organized groups can be analyzed by recording each attacker's *known accomplices*. We may assume that accomplices share knowledge of exploits, vulnerabilities, and assets [20,31]. 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 analyze a group with a common goal. The attributes of a *group* are:

- The group's membership: each identified member of the group should be analyzed individually, along with his role in the group. Also, leaders of the group should be identified.
- The group leaders' motives, and their goals and plans for the group.
- Interaction among group members, especially their cooperation and collective effectiveness.

Attackers need to be uniquely identified, in order to analyze them individually. Many of the attributes of attacker capability, as described earlier, can be used for identification. Additional identifying-attributes are the attacker's *peculiar work-habits*. For example, the attacker may habitually use an esoteric option for a particular shell command. An attacker may repeatedly display some illogical behavior, such as erasing the entire file system, except for the /tmp directory.

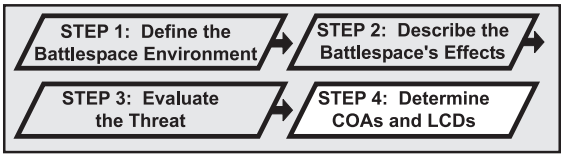

7. C-IPB's Step 4: Determine the threat's COAs and the LCDs

This step consists of two related tasks: determining the threat's COAs and identifying the LCDs.

Identifying LCDs, and their degree of likely compromise, is the ultimate objective of C-IPB. Determining the threat's COAs is a means to that end.

7.1. COAs

There are three types of attacker COAs: *possible*, *likely*, and *most-dangerous*. For C-IPB we are interested in possible and likely COAs. Estimates of the attacker's most-dangerous COAs are important for ARNC tactics and risk-management, but they are out-of-scope for C-IPB itself.

In the prior steps, we have collected the information needed for estimating COAs, namely the battlespace effects and the attacker's capabilities and intentions. Estimates of COAs will rarely be expressed as certainties. Gaps in knowledge of the battlespace-effects, and of the enemy, are natural and unavoidable. Gaps in knowledge can only be filled by assumptions about that which is typical and likely, based on prior experience and *basic intelligence*.

There are other difficulties in making COA estimates. Incorrect analysis in the prior C-IPB steps is possible – the attacker's deception-measures even seek to cause this. Also, the attacker may behave in a manner which is beyond our abilities to predict. Thought processes which are foreign to our own may be unknown and indiscernible. Capricious behavior, by definition, is not predictable (though consistent capricious behavior does reveal the attacker's planning process).

As mentioned earlier, the C-IPB process is dynamic. Each step of the process is updated as more is learned about the battlespace and the threat. If analysis is found to be incorrect, revisions will need to be made for it and for any conclusions drawn from it.

The accuracy of a COA estimate is a function of the uncertainties upon which it is based. Long-term estimates are typically based on more uncertainties than short-term estimates, so long-term estimates will tend to be more general, and less detailed, than short-term estimates. [15]

[15] In regard to long-term planning for battle, one of the most cited axioms in military-theory is, *no plan for battle survives the first encounter with the enemy* (attributed to Helmuth von Moltke, 1800–1891).

Estimates of the attacker's COAs can be for actions occurring during the short or long-term. The attacker's COAs will be based on his short and long-term plans, so estimates of the attacker's *planned COAs* are also needed.

The objective of C-IPB is to identify LCDs which are likely compromised in the past, present and future. So, estimates are needed for past, present and future COAs, both actual and planned.

7.2. LCD prediction

The objective of C-IPB is identifying LCDs, along with their *degree* of likely compromise. Identified LCDs can then be investigated for compromise, using established investigative techniques [9].

The purpose of determining the attacker's COAs is to identify LCDs. So, we are only interested in attacker COAs which can result in compromise, or which provide insight for identifying LCDs.

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. The ranking can be categorical (e.g., high, medium, and low) or ordinal. The LCD-set can include any device that has a vulnerability which can be compromised.

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 suspecting compromise (e.g., a particular vulnerability). Thus, a single device can have multiple entries in the LCD-set.

The following sections describe some principles and techniques for identifying COAs and LCDs.

7.3. Principles for estimating COAs

Estimates of the attacker's *possible COAs* can be obtained from analysis of the boundaries which

constrain his action. The constraints are the *battlespace effects* and the attacker's *capabilities*. (These constraints were analyzed in C-IPB's second and third steps.)

The attacker's *likely COAs* are a subset of his possible COAs. Likely COAs can be obtained from analysis of *how he acts* and of his *intentions*. How the attacker acts can be understood from his method of operation (MO) and from his access-patterns on the network. His access-patterns can reveal his intentions, plans and the time and location of his activity. (The attacker's MO and access-patterns were analyzed in C-IPB's third step.) Two other elements of intention are motives and appraisement. (They were analyzed in C-IPB's third step.) His choice of high-value-devices is especially important.

Note that possible COAs can be based largely on objective information, and likely COAs are based largely on estimative information.

7.4. Principles which govern COAs

7.4.1. Three basic actions of an attacker

There are three basic types of action which make up an attacker's COA: (1) attack, (2) gather intelligence, and (3) use network resources.

7.4.2. Simple predictions

Many of the attributes of the battlespace-effects, and the attacker's capability and intentions, can be used directly in the prediction of LCDs. Here are some examples:

- KCDs are likely to be used to attack other devices. Trust relationships can be exploited, as well as sensitive data obtained from the KCD.
- Exploits that the attacker has used are likely to be used by him again, if the opportunity arises.
- 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.

7.4.3. Economic principles

The Economics of Crime uses principles from economics as a means for understanding and predicting criminal behavior [30]. For our purposes, the economic attributes relevant to an attacker are: (1) his *valuation of network assets*, (2) his *costs* for exploitation of vulnerabilities, and (3) his *resources* for attacks.

Techniques for understanding these economic attributes, for the attacker, were described in earlier sections. The attacker's valuation of network assets was described in Section 6.5.2. The attacker's costs for exploiting vulnerabilities were described in Section 6.3.5. The attacker's perceived costs for potential losses were described in Section 6.5.1. The attacker's primary resources are described in Sections 6.3.1–6.3.4, 6.4 and 6.6. The section on MO describes the attacker's resource of time. An example of a personality trait which is a resource is boldness.

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

- When a device's vulnerability has an exploitation-cost that exceeds the attacker's resources, the device cannot be compromised. For example, a vulnerability's exploitation can require more skill than the attacker possesses.
- When a device's vulnerability has an exploitation-cost that is less than the attacker's resources, the device can possibly be compromised.
- When a device's vulnerability has an exploitation-cost that is less than the attacker's resources, and it is a highly valued-device, then the device is likely to compromised.
- Differences between asset-value and exploitation-cost can indicate the degree of likely compromise. An asset of little value, but of high cost to exploit, is not likely to be attacked. The converse holds as well.

- Devices with high asset-value or low exploitation-cost are likely to be compromised.
- Basic principles of economics are a guide for prediction: (1) Total costs cannot 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 costs.

7.4.4. 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 cannot predict the access (means) he will 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 LCDs. For example, the next attack target is likely to be a "weak link" (a device that has low exploitation cost) with attractive value.

7.5. Techniques for estimating COAs

Specification of COAs should include:

- what – the type of action, such as attack, gather intelligence, or use network resources,
- when – the time of the action,
- where – source, destination, and path taken,
- how – methods and techniques used,
- why – the objective sought.

Criteria for evaluating COAs are:

- suitability – the COA must have the potential for accomplishing the attacker's likely objective,
- feasibility – the attacker must have the ability and time to carry-out the COA,
- acceptability – the attacker must be willing to take the risks involved (more relevant to insider attacks),
- consistency – the COA must be consistent with the attacker's MO and recent activity.

COA estimates are hypotheses, and they need to be *verified*. When the COA includes attacks, the LCDs can be investigated for compromise, as mentioned earlier. Also, there may be points along the COA which offer opportunity for verification (e.g., traffic logs on a router). To verify future COAs, existing ID resources can be used, or new ones deployed. *Traffic logs can be used to determine the likelihood of past attacks.* For a device in the LCD-set, its network traffic can indicate the likelihood of a past attack, e.g., if there is no traffic from any of the KCDs to the LCD, this would lower the estimate of likely compromise.

8. Opportunities for future research

C-IPB is a proposed technique which is based on the authors' collective experience with IR, ID and IPB. This paper has presented: (1) the problem of locating LCDs during IR, (2) a process for locating LCDs, (3) the development of models of the battlespace and of the attacker, (4) the use of the Economics of Crime as a means for understanding and predicting an attacker's behavior.

There are opportunities for further developing C-IPB. C-IPB needs to be tested in actual IR cases. Also, data-recording and data-processing tools can be developed for handling the extensive amount of data collected for C-IPB. A fruitful source of ideas is the tools which the military has developed for its IPB process. (We are presently pursuing both of these research opportunities.)

Lastly, for C-IPB itself, there is a need for additional basic-intelligence on attacker capabilities, intentions and course-of-action.

C-IPB develops intelligence specifically for ID. Much of this intelligence could also be used for devising tactics for ARNC. For example, C-IPB presently determines the threat's possible and likely COAs. C-IPB could be extended to determine the threat's most-dangerous COAs. Also, more work can be done with C-IPB in discovering and exploiting the threat's vulnerability to detection.

There's benefit in collaboration, or alliances, among IR teams from around the world. CERT (http://www.cert.org) and FIRST (http://www.first.org) are some of the organizations established for this purpose. One means for collaboration is standardized IR processes. C-IPB could be used in developing a standardized ID process for IR.

Lastly, C-IPB could be adapted for use by law enforcement in criminal investigation.

9. Conclusion

For many networks, when a compromised device is discovered, it is likely to be the tip of the proverbial iceberg. There may be many other compromised devices. Locating all the compromised devices can be a very difficult task for the incident-responder.

A technique for identifying likely compromised-devices was presented, named *Cyber Intelligence-Preparation-of-the-Battlespace* (C-IPB). It is based on a military battlefield-intelligence process. C-IPB is summarized in Fig. 2. It provides a systematic process for building models of the battlespace, as well as the threat's capabilities, intentions, and courses-of-action. These models are used for identifying likely compromised devices.

Prior to an attack, an intrusion-detection system operates in anticipation of a general threat. Once an attack starts, the intrusion-detection system can deal less in the general and more in the particular – namely, particulars about attackers and attacked devices. C-IPB is an intrusion-detection technique which exploits the information that the attacker reveals about himself and about network vulnerabilities.

In the contest between attacker and incident-responder, which combatant will have superior battlespace intelligence?

Acknowledgements

- Funding for this work was provided by DARPA/ISO as part of the Information Assurance and Survivability programs under federal contract F30602-96-C-0325.
- We would like to thank five key-contributors whose assistance made this cross-discipline research possible:
  - Dr. Christopher Bassford, Professor of Strategy, National War College, DOD.
  - Colonel G.I. Wilson, I Marine Expeditionary Force, USMC.
  - Dr. E.C. Pasour, Professor of Agricultural Economics, North Carolina State University.
  - Major Andy F. Knights, Wolfpack Battalion, Army ROTC, North Carolina State University.
  - The Ludwig von Mises Institute, http://www.mises.org, Auburn, Alabama.
- Very helpful advice was received from: Dr. Maurice Godwin, Criminal Investigative Psychologist, University of Alaska Anchorage; Dr. Gary Weinberg, College of Engineering, NCSU; Gunnery Sergeant Woody Biggs, I MEF, USMC; Lt. Colonel Neil Garra, US Army; Major Jerry Schlabach, US Army.

Glossary

AA	avenue of approach
AO	areas of operation
AOI	areas of interest
ARNC	attack repair, neutralization and containment
C-IPB	Cyber-IPB
CD	compromised device
COA	course of action
ID	intrusion detection
IDS	intrusion–detection system
IPB	intelligence preparation of the battlespace
IR	incident response

KCD	known compromised device
LCD	likely compromised device
MO	method of operations
USMC	US Marine Corps

References

[1] E. Amoroso, Intrusion Detection, Intrusion.Net Books, 1999.

[2] Maximum Security: A Hacker's Guide to Protecting Your Internet Site and Network, Sams.net Publishing, 1997.

[3] T. Aslam et al., Use of a taxonomy of security faults, COAST Laboratory, Technical Report TR-96-051, Purdue University, 1996.

[4] J. Cameron, Method in Software Development, JSP & JSD: The Jackson Approach to Software Development, IEEE Computer Society, Silver Spring, MD, 1983.

[5] C. von Clausewitz, On War, Princeton University Press, Princeton, NJ, 1832.

[6] D. Denning, Cryptography and Data Security, Addison-Wesley, Reading, MA, 1982.

[7] DilDog, Back Orifice 2000, http://www.bo2k.com/, 2000.

[8] D. Farmer et al., Improving the security of your site by breaking into it, full text at http://www.cerias.purdue.edu, comp.security.unix, December 1993.

[9] R. Firth et al., Detecting signs of intrusion, Full text at http://www.cert.org, Carnegie Mellon University, Software Engineering Institute, Security Improvement Module CMU/SEI-SIM-001, 1997.

[10] US Army Intelligence Center, FM 34-130 Intelligence Preparation of the Battlefield, Full-text at http://155.217.58.58/atdls.htm, US Army, 1994.

[11] R. Forno et al., The art of information warfare: insight into the knowledge warrior philosophy, Full text at http://www.upublish.com, Upublish.com, 1999.

[12] D. Freedman et al., At Large: The Strange Case of the World's Biggest Internet Invasion, Simon & Schuster, New York, 1997.

[13] Fyodor, nmap, http://www.insecure.org/, 2000.

[14] R. Heuer, Psychology of intelligence analysis, Full text at http://www.odci.gov/csi/, CIA, Center for the Study of Intelligence, 1999.

[15] D. Icove, Collaring the cybercrook: an investigator's view, IEEE Spectrum, 1997.

[16] Invisible Evil, Hacking Kit v2.0.b, distributed on the Internet, 1997.

[17] G.M. Godwin, Hunting Serial Predators, CRC Press, Boca Raton, 2000.

[18] A.A. Jalali et al., The other side of the mountain: Mujahideen tactics in the Soviet–Afghan war, US Marine Corps, Studies and Analysis Division, 1999.

[19] G. Kim et al., The design of a system integrity monitor: Tripwire, Full text at http://www.cerias.purdue.edu/, COAST TR 93-01, Department of Computer Sciences, Purdue University, 1993.

[20] K.P. Kossakowski et al., Responding to intrusions, Full text at http://www.cert.org, Carnegie Mellon University, Software Engineering Institute, Security Improvement Module CMU/SEI-SIM-006, 1999.

[21] U. Lindqvist et al., How to systematically classify computer security intrusions, in: Proceedings of the 1997 IEEE Symposium on Security and Privacy, IEEE Computer Society Press, Silver Spring, MD, 1997.

[22] S. McClure et al., Hacking Exposed: Network Security Secrets and Solutions, Osborne/McGraw-Hill, New York, 1999.

[23] Doctrine Division, MCDP 1-3 Tactics, Full text at http://www.doctrine.usmc.mil, US Marine Corps, MCCDC, 1997.

[24] Doctrine Division, MCDP 2 Intelligence, Full text at http://www.doctrine.usmc.mil, US Marine Corps, MCCDC, 1997.

[25] MCI 7510B Tactical Fundamentals, US Marine Corps, Marine Corps Institute, 1984.

[26] Doctrine Division, MCWP 2-1 Intelligence Operations, Full text at http://www.doctrine.usmc.mil, US Marine Corps, MCCDC, 1998.

[27] C4I & MCIA, MCWP 2-12 MAGTF Intelligence Analysis and Production, Full text at http://www.doctrine.usmc.mil, US Marine Corps, MCCDC, Doctrine Division, 1999.

[28] L. von Mises, Human Action, first ed., Ludwig von Mises Institute, 1949.

[29] C. O'Hara, Fundamentals of Criminal Investigation, Thomas, Springfield, 1994.

[30] M. Reynolds, Crime by choice: An economic analysis, Fisher Institute, 1985.

[31] S. Romig, State of the hack, Full text at ftp://ftp.net.ohio-state.edu:/users/romig/talks/state-of-the-hack, Ohio State University, UTS Network Security Group.

[32] A. Rubin et al., Web Security Sourcebook, Wiley, New York, 1997.

[33] E.E. Schultz et al., Responding to computer security incidents: guidelines for incident handling, Full text at ftp://ciac.llnl.gov/pub/ciac/ciacdocs/ihg.txt, Technical report from Department of Energy, LLNL, 1990.

[34] M. Slatalla, Masters of Deception: The Gang That Ruled Cyberspace, HarperCollins, 1995.

[35] Super et al., The Undernet Botdocs: An Introduction to IRC Bots, http://www.undernet.org/documents/, Internet document, circa 1997.

[36] West-Brown et al., Handbook for Computer Security Incident Response Teams (CSIRTS), Full text at http://www.sei.cmu.edu/, Carnegie Mellon, Software Engineering Institute, Handbook CMU/SEI-98-HB-001, 1998.

Jim Yuill, Ph.D. candidate, Computer Science Department, North Carolina State University, jimyuill@pobox.com, http://www.pobox.com/~jimyuill

Dr. Felix Wu, Assistant Professor, Computer Science Department, North Carolina State University, wu@adm.csc.ncsu.edu, http://www.csc.ncsu.edu/directory/bio/FWu.html

Jim Settle, Chief, FBI Computer Crime Squad (Retd.); presently, President, Settle Services in Technology, settle@netcom.com

Dr. Feming Gong, Director, Advanced Networking Research, MCNC, gong@mcnc.org, http://www.anr.mcnc.org/staff-pages/FengminGong.html

Rick Forno, Director of Security, Network Solutions, Inc., rforno@ibm.net, http://www.infowarrior.org

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

John Asbery, President, Marine Corps Intelligence Association; Master Gunnery Sergeant, USMC (Retd.); presently at The Boeing Company, johnny.n.asbery@boeing.com