Tactical Operations and Strategic Intelligence:
Sensor Purpose and Placement

TR-2002-04.02

9 Sep 2002

Todd Heberlein
Net Squared, Inc.
todd@netsq.com

This paper is motivated by discussions with a number of people over where to place network-based intrusion detection sensors. The answer depended what you wanted to do with the sensor information. This lead to an examination of three types of activities that use intrusion detection sensors: (1) tactical operations, (2) aggregated tactical operations, and (3) strategic intelligence. We examine several steps that a site can take to optimize a sensor's use in tactical operations. We briefly look at the outsourcing of tactical operations; although, we limit the amount of discussion because we believe the value of such activities is limited. Finally, we look at using sensor grids to develop strategic intelligence. As this role is perhaps the least understood of all the uses of intrusion detection sensors, we spend considerable time fleshing out some of the concepts.

Table of Contents

List of Figures

1 Introduction

Figure 1 shows a common network design. The design partitions the world into three areas: (1) the outside world, (2) a demilitarized zone, and (3) a protected network. The demilitarized zone (DMZ), technically part of an organization's network, typically contains servers that need to be accessed by the general public, such as public web and FTP servers. The rest of an organization's network, which generally does not need to be publicly accessible, is protected behind a firewall. Many sites, both commercial and military, deploy their network-based intrusion detection sensors in the DMZ.

Unfortunately, sensors placed in the DMZ also generate huge numbers of alerts, often numbering in the thousands or tens of thousands of alerts per day. For example, IBM's Real-timed Intrusion Detection Service (RTIDS), for a sample of 27 Cisco intrusion detection sensors at customer sites for a one-month period, received on average over 14,000 alerts per sensor per day [Mang 99]. Similarly, an analysis of two months of Snort Alerts from the Air Force's Rome Labs generated a median of over 10,000 alerts per day and an average of 59,000 alerts per day. The operators of the Rome Snort sensor have told us that after significant tuning they have reduced the alerts to roughly 4,000 per day. Finally, DARPA program managers frequently cite similar stories of overwhelming data flows from intrusion detection sensors as a reason why DARPA must develop new technologies.

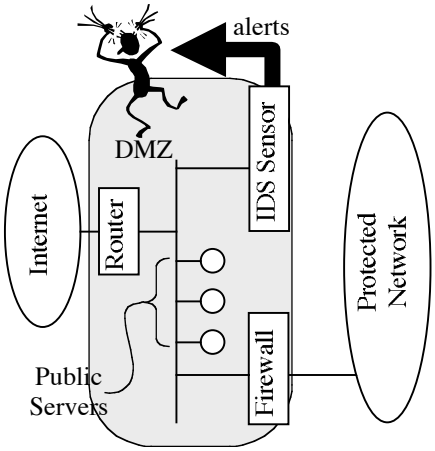

Figure 1: Classical Sensor Placement

If an organization's mission includes providing an operational intrusion response capability to reported attacks, then the environments just described create a nearly impossible situation. For example, if the organization's service level agreement (SLA) to its customers requires it to manually sign-off on every reported attack message from a managed sensor, then even at the low-end of 4,000 alerts per day a security administrator would have on average 21 seconds to determine if a reported attack needs further investigation. If the reported attack does need further investigation, someone else would probably need to perform this duty because many more alerts that need to be triaged will be coming down the pipeline – on average a new alert every 21 seconds, continuously 24 hours a day, 7 days a week.

Some intrusion detection management systems provide statistical roll-up capabilities to group many alerts into a single report for display purposes; however, this does not necessarily solve the fundamental problem. If host *bad_guy* launches an RPC attack *xdr_overflow* against 16,000 hosts within a site, an intrusion detection management station may instead of displaying the 16,000 alerts from the intrusion detection sensor display only a single report that states 16,000

machines were attacked. However, if the organization's mission includes attack response/management, then security administrators still need to dive into those 16,000 alerts to determine which if any may have been successful and then decide what to do about those that did succeed.

In a report last year titled "Before Applying New Technologies" [Hebe 01], we argued that much of the problem of overwhelming sensor alerts could be resolved through better procedures, processes, and engineering, and that these steps should be considered before investing heavily in new technologies. Sections 2.1 and 2.2 present two of the suggestions from that report.

But we also recognize that the role of intrusion detection is changing. Today there is a greater emphasis on creating large-scale sensor grids, where sensor reports from many sites are coordinated for some purpose. The role of the sensors at a site will vary depending on what the purpose is for forming the sensor grid.

Figure 2 shows a spectrum of purposes for which a sensor may be used. On the far left is the single site sensor configuration. In this role the sensor's primary purpose is to help analysts manage attacks. Typically the analysts prefer a minimum number of attack alerts because each one needs human attention. On the far right is the threat intelligence sensor grid. For this purpose, the analysts prefer a rich sensor feed because the reports will be processed through statistical algorithms and not be processed one-by-one by humans. And in the middle is a hybrid effort, forming a sensor grid from many sites but where analysts are still responsible for managing each attack.

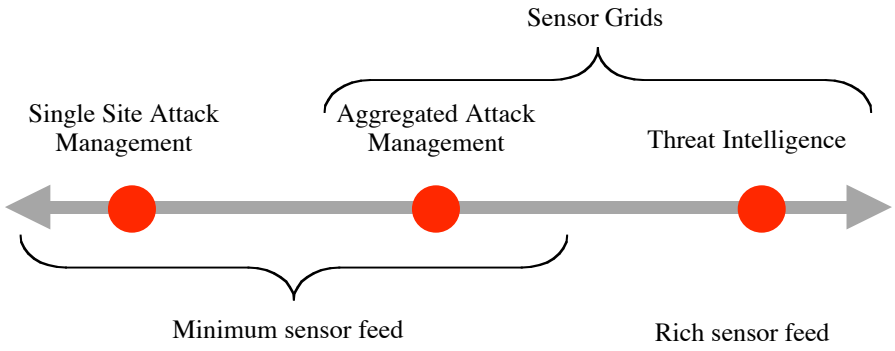

Figure 2: Sensor Purpose Spectrum

In this report we look at each of these roles. In Section 2, Tactical Operations, we look at sensors used in single site attack management. If we use war a analogy, this is where the ground operations take place, and where combat is personal. The terrain must be prepared, and attacks that make it past the initial defenses must be managed one-by-one. In this section we discuss some simple steps to optimize these tactical operations.

In Section 3, Aggregated Tactical Operations, we look at a common business model where attack management is outsourced to a central organization. As in the previous section, in this role individual attacks must still be managed, the difference is that the management is taking place at a remote location.

In Section 4, Strategic Intelligence, we look another purpose for a sensor grid: to generate a detailed understanding of the actual threats that individual sites and the network as a whole are facing. Unlike the previous two roles, here we are not concerned with managing specific attacks, Similarly, while the other two roles prefer to minimize sensor reports to only those reports that require human attention, strategic intelligence prefers a rich sensor feed in order to divine a

deeper understanding of the attacks. Many organizations try to combine aggregated tactical operations with strategic intelligence, but because their data needs and end products are so different, we prefer to treat them as separate activities. Throughout this section we examine some of the many ways a strategic intelligence organization can create value.

2 Tactical Operations

Tactical operations is concerned with managing individual attacks against a site as well as preparing a site prior to an attack. It essentially encompasses all actions by security, network, and system administrators that change the state of the network with respect to security. This can include installing new services, applying patches, modifying configuration files on applications operating systems, and firewalls, and blocking, investigating, and cleaning up after specific attacks.

In the early days of deployed intrusion detection systems (circa 1991) management and security administrators often wanted to know about, and in many cases investigate, every attack against their site whether it was successful or not. Today, with sites under continuous attacks all day every day, hands-on response to each attack instance is nearly impossible.

Instead, today's tactical response to attacks should focus only on attacks that need human attention, and these are primarily attacks that are, or may be, successful. Attacks that fail because the targeted service does not exist, the attack service is patched, or the attack is stopped by other mechanisms such as firewalls do not need to be, and should not be, handled directly by humans. This section discusses some approaches to help the security administrator focus only on the potentially successful attacks.

2.1 Moving the Primary Sensor

The simplest and probably the most effective step to reducing the volumes of reports analysts must process is placing the sensor behind the firewall (see Figure 3). The firewall should be the first line of defense, not the second. An effectively configured and managed firewall (or set of firewalls) should screen out most attacks a site could potentially face. Because the sensor will never see these attacks, this should substantially reduce the number of reports an analyst must examine.

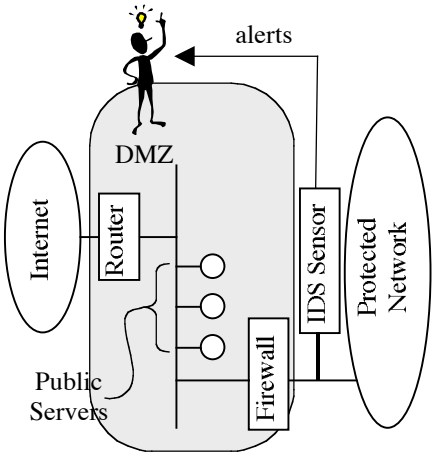

Figure 3: Alternative Sensor Placement

There are at least three common reasons why organizations continue to place IDS sensors in front of a sites firewall. First, there are machines inside the DMZ that need to be protected,

and placing the sensor behind the firewall hides attacks against those systems from sensor. Because of this, an organization may want to use two IDS sensors: the primary one placed behind the firewall that detects attacks against most of site's computers and one in the DMZ tuned to only look for attacks against the systems in the DMZ. Since the DMZ typically only includes a handful of machines, the sensor tuned to protecting them can be relatively small, and the number of reports generated by it should be relatively light.

A second reason that a primary intrusion detection sensor remains in the DMZ instead of behind a firewall is that the customer site does not trust the organization performing the intrusion detection service enough to allow the organization to access systems behind the firewall. This is most often the case when a third-party such as intrusion detection monitoring service such as IBM's Real-Time Intrusion Detection Service or the Air Force Information Warfare Center (AFIWC) is monitoring customer sites. This is largely a political issue and not a technical one, but if the organization is tasked with responding to attacks, keeping the primary sensor in the DMZ will severely hurt the organization's ability to perform its job.

A third reason that the primary sensor remains in the DMZ is that the monitoring organization wants to know about all attacks against a site. There are at least two reasons that knowing about all attacks, even if the vast majority of attacks are blocked by the firewall, can be valuable. Knowing about all attacks provides a clearer view of the overall threat picture, and we cover this topic in more detail in Section 4.

Another reason for knowing about all attacks is that it can help in elevating the warning of the few attacks that make it pass the firewall. For example, if host *bad_guy* attacks 1,000 computers at a site, of which 999 are blocked by the firewall, the sensor inside the firewall would only see a single connection from *bad_guy*. The sensor (or the security administrator reviewing an output from the internal IDS sensor) might be more inclined elevate the threat level posed by that one connection if it (or he) was aware of the other 999 connection attempts.

In either case, however, the attack attempts that are prevented from reaching the protected network by the firewall should be treated differently than those that actually make it past the firewall. Alerts associated with activity blocked by the firewall should not be considered actionable messages that should handled by tactical response teams.

2.2 Integrating Vulnerability Knowledge

After placing the primary IDS sensor behind the firewall, the second most effective step in reducing the volume of alerts generated by the sensor is integrating vulnerability information into the analysis. Essentially vulnerability information is treated as a filter, removing many alerts before presenting them to a security administrator (see Figure 4). For example, if host *bad_guy* launches an attack against a Microsoft IIS vulnerability on an Apache web server running on a Sun server, the attack will certainly fail, so the alert does not necessarily need to be sent to a tactical response team.

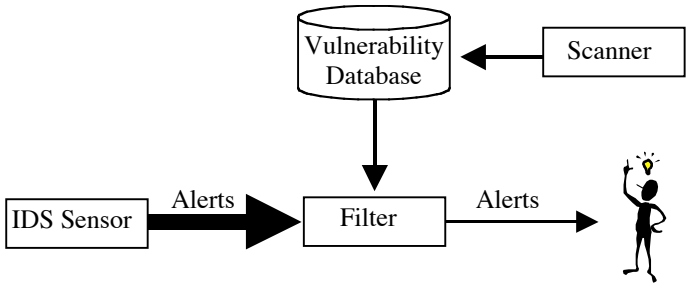

Figure 4: Integrating Vulnerability Information

Technically the greatest challenge is to clearly relate vulnerabilities identified by the scanner to reports generated by the sensor. Ideally both systems would use a common identification system. For example, when a sensor detects a particular attack against a web server, ideally the report would also include an ID (e.g., a CVE identifier) identifying the vulnerability that the attack exploits. Many different attacks might exploit the same vulnerability (one may propagate a worm while another may fork a shell), so it is important to distinguish between vulnerabilities and attack signatures. With this capability, a filter could look into the database to determine if the system targeted in the attack was tested for that vulnerability, and if it was tested, whether the system was vulnerable. If the system was tested for the vulnerability and was secure, then the report does not have to be submitted directly to the analyst. On the other hand, if the vulnerability was tested and the system was vulnerable, then the attack report should be elevated to a higher level. Attacks for which success or failure cannot be determined should also be forwarded to the analyst.

3 Aggregated Tactical Operations

Aggregated tactical operations bring many of the activities mentioned in Section 2 for many sites under a single organization. The Air Force pioneered aggregated tactical operations with respect to intrusion detection monitoring and response with the build up of their Automated Security Incident Measurement (ASIM) program in the early to mid 1990s. Later, one of the ASIM chief architects, Dan Teal, left the Air Force and formed WheelGroup to provide similar services commercially. Since then, a number of organizations including ISS, Counterpane, and IBM, have all offered aggregated tactical operations for various types of security capabilities.

Since many of activities that fall under the umbrella of tactical operations require system administration privileges on local machines (e.g., to apply patches), aggregated tactical operations tend to focus on managing security devices such as network-based intrusion detection sensors and firewalls (e.g., keeping signatures and configuration files up to date), what Counterpane's Bruce Schneier calls Managed Security Services (MSS), and monitoring the output from these security devices to detect attacks, what Schneier calls Managed Security Monitoring (MSM) [Schn 01].

Typical arguments in favor of aggregating tactical operations include cost benefits from economies of scale, a deeper pool of experts to respond to unusual threats, and better intelligence based on more comprehensive view of the threat environment. The first argument, economies of scale, is truly part of tactical operations (e.g., requires hands-on response to specific attacks), but the latter arguments fall more closely under the category we call strategic intelligence, which we cover extensively in Section 4.

Unfortunately, we believe the financial benefits from economies of scales provided by Counterpane's Bruce Schneier to justify their business model [Schn 01] is based on somewhat flawed assumptions. In "Managed Security Monitoring: Network Security for the 21st Centruy" Schneier argues that for a site to provide its own security monitoring 24 hours a day and 365 days a year, the site would require at least five full-time employees. Add in supervisors and personnel with special security skills, and the cost of Counterpane's service suddenly looks attractive.

The flaw in the argument is that few sites hire individuals just for tactical security operations, rather these functions are usually part of a network or system administrator's normal job responsibilities. Hiring a managed security monitoring service such as Counterpane probably will not result in the ability to reduce a half dozen to a dozen system administrators from the payroll, so the cost savings argument is not as strong as it first appeared.

Also, tactical response often requires detailed knowledge of the actual site, including policies, an organization's missions and priorities, physical and logical topologies and

dependencies in the network, and personnel matters (e.g., new hires and recent dismissals). Local network and system administrators are more likely to be up to date on these issues than people an outsourced monitoring operation.

To summarize this section, while a centralized security organization has value (see Section 4), much of the hands-on tactical operation needed for security (updating operating systems, applying patches, understanding a user's unusual behavior, and deciding whether or not to block access by specific user or IP address) can often be done best by local security, network, and system administrators than by personnel at a remote centralized organization.

4 Strategic Intelligence

Strategic intelligence is concerned with understanding the threat environment in order to optimize tactical operations. This role is similar to the mission of the United States Center for Disease Control and Prevention's (CDC) Epidemiology Program Office (EPI)[1] and National Center for Infectious Diseases (NCID)[2] as well as the World Health Organization's (WHO) Communicable Disease Surveillance and Response (CSR) group[3].

Whereas tactical operations focuses on individual attacks, strategic intelligence focuses on the overall threat picture. There may be 500 attacks going on at any moment in time, but a strategic intelligence group should not necessarily concern itself about any of these individual attacks. A strategic intelligence group should focus on the nature of those 500 attacks. How similar or different are today's attacks from those of yesterday's attacks, or last month's attacks? Is there a new threat? Is a new threat expanding? Is an old threat receding?

In addition to looking at global-scale issues raised in the previous set of questions, a strategic intelligence group also looks at issues effecting individual sites that can only be gleaned from examining data from a global perspective. For example, what attacks are likely to hit a site in the future? Is there an attack that is too subtle to be noticed at an individual site but that may pose an important threat?

The answers to all these questions, provided by interpretation of analyses of data from a global sensor grid, are bundled into actionable information and delivered to the appropriate people (e.g., network and system administrators in the field) in a timely fashion. In the end, only the field operators responsible for local systems and security can affect change, and it is the role of the strategic intelligence group to make these people as effective as possible.

This process provided by a strategic intelligence group is captured succinctly by a definition of "surveillance" provided by WHO and shown in Figure 5.

> Surveillance is the *ongoing systematic* collection, collation, analysis and interpretation of data; and the dissemination of information to those who need to know in order that action may be taken.

Figure 5: Disease Surveillance [WHO 99]

[1] http://www.cdc.gov/epo/
[2] http://www.cdc.gov/ncidod/
[3] http://www.who.int/emc/

In addition to providing actionable information to individual security and system administrators, a strategic intelligence group should provide expertise to individual sites to help them understand the nature of a potential new threat. No individual site can field enough expertise in enough areas to comprehend and develop countermeasures for all possible threats they may face, so part of the role of a strategic intelligence group is to provide backup experts to help these sites.

Sections 4.1 through 4.3 examine some of the unique capabilities that a strategic intelligence organization can provide.

4.1 Attack Prediction

In an ideal world, all systems would be patched as soon as patches are available, but real-world networks rarely approach this ideal situation. In a typical week in 2000, Bruce Schneier noted 13 vulnerabilities reported [Schn 00]. During the same week in 2001 Schneier counted 19 patches released [Schn 01]. In 2001, the Computer Emergency Response Team (CERT) identified 2,437 vulnerabilities, an average of 47 new vulnerabilities each week [Schw 02]. Tracking and patching all systems at a moderate sized site, especially when many different administrators, including the users themselves, manage those systems can quickly become a nightmare.

In part, this is the need that signature-based intrusion detection systems fill. Since most intrusion detection systems primarily detect known attacks against known vulnerabilities, (for which there already exists a patch or countermeasure), their primary use by tactical response teams is to identify attacks against systems that have not been patched (see Section 2.2).

Unfortunately, cleaning up after a successful attack into a system is typically at least an order of magnitude more expensive than securing the system before the attack. Since a site is usually unable to patch all systems all the time, an ideal situation would be for the site to know which attacks would be launched against their sites so the site can focus its preparations on defending against those attacks. Fortunately, of the dozens of vulnerabilities reported each week, only a small number are actually exploited in attacks. For example, of the 2,437 vulnerabilities identified by CERT in 2001, less than 1% of them were exploited in actual attacks [Schw 02]. The secret to success is for a site to know which 1% will be exploited before they are attacked, and this is where strategic intelligence plays an important role.

Figure 6 illustrates the current model. A site's sensor boundary only extends to the site's edge, so the site can only detect attacks that are immediately upon it. In a sense, current architectures create a very myopic view of the threat environment.

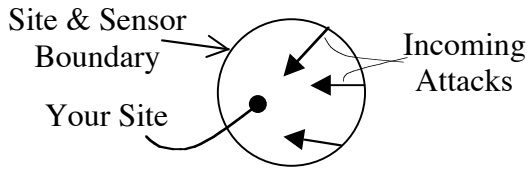

Figure 6: Myopic Sensor View

However, by integrating sensor data from many sites a strategic intelligence organization can calculate a reasonable probability of an individual site encountering a threat before that threat occurs (see Figure 7). In effect, we extend the sensor capabilities beyond the range of an individual site's boundary (i.e., "over-the-horizon" intrusion detection) to perform early detection. With early warning, a site can prepare for the threat as opposed to reacting to it when it occurs.

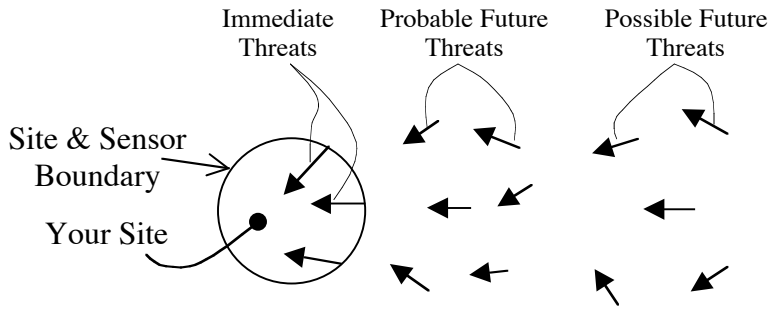

Figure 7: Over-the-Horizon Threat Detection

The next two sections discuss techniques that a strategic intelligence organization can use to help individual sites determine probable and possible future threats so that the sites can prepare for them.

4.1.1 Global Trends

The first and simplest approach to predicting attacks that individual sites will see is to look at the attacks that are most prevalent on the Internet (or the portion of the network monitored by the strategic intelligence organization). This essentially creates a "top 10" list of attacks that a site should prepare for. In other words, if you only have time to patch 10 vulnerabilities, these are the 10 you should patch. In practice, more than just 10 should be available. This approach resembles the SANS Incidents.org InternetStormCenter's "Top 10 Ports" list[4] and Amazon's "Top 100 Bestsellers" list[5].

In addition to the most active vulnerabilities being targeted, a strategic intelligence organization can provide trending information to sites. Trending indicates which attacks that have a relatively low ranking today will probably have a higher ranking over the next several days or weeks. The easiest approach to providing this information is to compare an attack's most recent activity (e.g., last five days) to its longer-term activity (e.g., the previous 30 days). If the number is positive, it is trending up; if it is negative, it is trending down. Examples of such measurements include the SANS InternetStormCenter's Trends page[6] and Amazon's "Movers & Shakers"[7].

4.1.2 Victim Profiles

The global trends approach is a generic approach to predicting attacks that *any* site might be likely to see. The only important factor is that the strategic intelligence organization collects attacks from a large number of sites. If an attack has already hit 20% of the sites, the other 80% of the sites should be notified so they can prepare for it. However, no one site is treated different from another site.

"Victim profiles" provides an approach that can further refine the prediction for a specific site. The underlying model is that some sites will have similar attack patterns sent against them, and identifying clusters of similar victims is useful for refining predictions. If site *widgets.com* is in victim group **A**, and 40% of sites in group **A** have observed attack *xdr_overflow* even though

[4] http://isc.incidents.org/top10.html
[5] http://www.amazon.com/exec/obidos/tg/browse/-/549066/hot/1/103-7992537-9327010
[6] http://isc.incidents.org/trends.html
[7] http://www.amazon.com/exec/obidos/tg/new-for-you/movers-and-shakers/-/books/ref=pd_gw_msgr/103-7992537-9327010

only 2% of the global population has seen the same attack, then *widgets.com* should be warned about a strong potential of seeing attack *xdr_overflow* soon.

Figure 8 shows one approach to identifying victim profiles. A site is represented by large vector representing its history of attacks. For example, the vector may contain counts of 1000 attack types that the strategic intelligence organization is tracking, the IP addresses of 500 recent source networks that have launched an attack against a site, and the top periods of the day that a site sees attacks. A distance function $f(\mathbf{V}_1,\mathbf{V}_2)$ that measures a distance between two site vectors is supplied to a clustering algorithm, and the resulting clusters are the victim groups. This approach is similar to how Amazon.com generates custom recommendations for its customers.

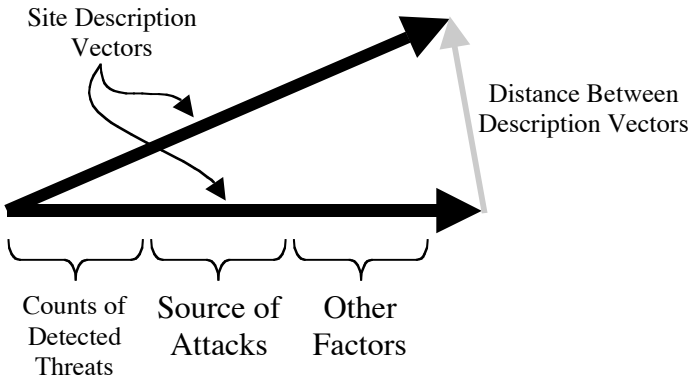

Figure 8: Measuring Site Similarity

4.2 Identifying Important Attacks

There is a difference between the most actives attacks a site, or groups of sites, might see and the most important attacks a site might experience. Many attacks a site might see can be considered random acts of violence. These include the mindless Code Red and Nimda attacks that continue on the Internet as well as popular scripts run by amateurs (e.g., script kiddies) who thought they would just try the tool against a site. In raw numbers, these attacks often dominate the "top 10" or "top 20" attacks a particular site will see. However, the most important attacks are the ones that will most likely cause the most damage (however that is calculated), and these attacks may be numerically well down on the attack list. In particular, these may be attacks developed by highly motivated attackers (and perhaps sponsored attackers) with very specific goals in mind.

For example, power grid sites are probably bombarded by huge numbers of the same random acts of violence attacks that every other site is. However, if there is some specific attack that is relatively unique to power grid sites, they may represent a concerted effort to take down the nation's power grid, so these attacks should be given special attention despite being numerically insignificant.

Amazon.com accomplishes something similar through a technique they call "Purchase Circles"[8]. Figure 9 illustrates this with book purchases. The left column shows the top four selling books to Air Force customers. This list is very similar to the overall best-selling list for Amazon.com, and this would be like our random acts of violence attacks mentioned previously. The right column, on the other hand, shows the top four selling books that are relatively unique to

[8] http://www.amazon.com/exec/obidos/subst/community/community.html/ref=gw_hp_ls_1_10/103-7992537-9327010

Air Force customers. It provides a clearer picture of what makes the Air Force unique from other organizations. We should be able to provide such a unique view for organizations based on their attack profiles.

Books ranked highest in sales overall at the Air Force	Books uniquely popular to the Air Force
(1) Harry Potter and the Goblet of Fire (Book 4)	(1) Air Power: A Centennial Appraisal
(2) Body for Life: 12 Weeks to Mental and Physical Strength	(2) Writing Guide for Air Force Efficiency Reports
(3) Harry Potter and the Sorcerer's Stone (Book 1)	(3) They Also Flew: The Enlisted Pilot Legacy, 1912-1942
(4) Who Moved My Cheese? An Amazing Way to Deal with Change in Your Work and in Your Life	(4) Joint Air Operations: Pursuit of Unity in Command and Control, 1942-1991 (A Rand Research Study)

Figure 9: Best Sellers vs. Uniquely Popular

4.3 Detecting New Threats

Through roughly the first decade of intrusion detection research and development, the various efforts largely steered clear of signature-based intrusion detection [Mukh 94]. Signature-based detection schemes were seen as ad hoc, where a new solution had to be hand-crafted for each new attack. Signature-based detection systems were also faulted for their inability to detect new attacks, or even variations on existing attacks.

While these criticisms certainly carried a certain amount of truth, beginning in the mid 1990s signature-based intrusion detections began to dominate the market of deployed intrusion detection systems. Thus, in the end there is a large installed based of intrusion detection systems, but they are relatively ineffective at detecting new attacks.

We believe the primary reason for signature-based systems domination of the intrusion detection market is that in most tactical operations (see Section 2), schemes that might detect new attacks (generally referred to as anomaly-based intrusion detection systems) (1) generate too many false alarms, and (2) do not provide actionable information. However, for strategic intelligence organizations, these issues are generally not a problem, so we believe that strategic intelligence organizations may actually provide a marketplace for anomaly intrusion detection systems.

In Section 4.3.1 we look into the problem of anomaly detection for tactical operations. In Section 4.3.2 we briefly discuss a limitation in antivirus software and how strategic intelligence operations may address this problem. In Section 4.3.3 we look into detail of how anomaly detection can be used to identify subtle new attacks.

4.3.1 Problem with Anomaly Detection

As mentioned in Section 2, personnel in tactical operations generally want to be notified of an attack only when the attack requires their attention. Furthermore, in an ideal situation, an

intrusion detection system should not just report an attack but also provide the user with actionable information, and generally this is fairly doable.

When a person creates a new signature for an attack, they are usually aware of (1) the attack that the signature should detect, and (2) the vulnerability the attack exploits. In systems like Snort it is very easy to include both the attack name and an ID for the vulnerability it exploits (e.g., a CVE identifier[9]) with each report of a detected attack. Likewise, when a site scanner looks for a vulnerability, it should be able to provide a well known ID for each vulnerability it has found. By combining the attack report, analysis from the vulnerability scanner, and a service such as ICAT[10] that links a CVE ID to a set of patches and links for additional details, an intrusion detection system can easily generate a report such as the one in Figure 10, column A.

	A		B
Target:	128.131.7.2 : 161	**Target:**	128.131.7.2 : 161
Attacker:	128.120.56.31 : 5611	**Attacker:**	128.120.56.31 : 5611
Attack Name:	xdr_router_crash	**Attack Name:**	unknown
Vulnerability ID:	CVE-2002-0391	**Vulnerability ID:**	unknown
Vulnerable:	Yes	**Vulnerable:**	unknown
Damage:	Crashes Cisco routers	**Damage:**	unknown
Link to Patch:	Cisco_patch	**Link to Patch:**	none
Details:	Security Focus CERT CC	**Details:**	none

Figure 10: Signature vs. Anomaly Reporting

The report in column A is what we call *actionable information*: it tells you (1) what the attack was, (2) whether you were vulnerable and need to do something, (3) where to get a patch to secure the system, and (4) where you can go for additional details.

An anomaly-based intrusion detection system, on the other hand, is more likely to generate a report that resembles Figure 10, column B. It might detect something suspicious, but it cannot give you a name for the attack, it cannot tell you about a specific vulnerability that needs to be addressed, and it cannot tell you what you need to do about it. This is definitely not actionable information.

4.3.2 Virus Detection: Yes, No, Maybe

While most people think virus detectors simply look for strings in files, antivirus software developers have developed a number of techniques to combat the ever-evolving forms of deception approaches that virus writers have deployed. Unfortunately, virus writers' techniques have become "so effective that many mainstream antivirus products are still unable to detect such infections months after the code's release" [Nach 02a].

Carey Nachenberg, chief researcher at the antivirus company Symantec, has said that they have developed algorithms that use heuristics that can identify that a file *might* contain malicious code, but Symantec did not think that such an approach would work in the marketplace. Consumers want a definitive answer, *yes* or *no*, as to whether the file contains malicious code. The consumer must make a decision about whether to open the file, and a recommendation of

[9] http://cve.mitre.org/
[10] http://icat.nist.gov/icat.cfm

"maybe" is generally not an acceptable answer. Thus, in the end, these detection techniques are generally not deployed, or, when they are deployed, they are not activated by default. [Nach 02b]

A strategic intelligence organization, however, could make use of "maybe" answer from antivirus software. For example, suppose Symantec signs a site license with the Air Force to deploy their antivirus software on all email hubs and web and FTP proxies. The antivirus software is configured to block files that are definitely infected, allow all files that are definitely not infected to flow through, and allow potentially infected files (the "maybes") to flow through depending on the locally active policy. However, all "maybe" files are also shipped, probably in an encrypted form, to an Air Force strategic intelligence unit, where the analysts are cleared for highly sensitive data (e.g., they should be able to analyze "maybe" infected files that are marked "top secret"). The expert analysts can examine the file in more details, and they can compare files against other "maybe" files observed throughout the network. If very stealthy malicious code that cannot be definitively identified in isolation as malicious is spreading, the Air Force strategic intelligence organization will observe an increased trend of "maybe" files being flagged. That is, one "maybe" might not be suspicious, but 30 "maybe" files should raise concern.

This approach, which is very specific to antivirus software, is similar to what we discuss in more detail in the next section, Section 4.3.3. The point is, that a strategic intelligence unit could create a viable marketplace into which the antivirus software vendors could sell their software. The typical consumer will not accept detection software that answers *yes*, *no*, or *maybe*. However, a strategic intelligence unit can take advantage of the *maybes*.

4.3.3 Detecting New Subtle Systematic Activity

In this section, we illustrate an approach to detect subtle activity that is spread across many sites. We begin by describing a simulation that we created to illustrate the issues. Next, we look at how a single site could eventually detect a very subtle attack. Finally, we look at how a strategic intelligence organization could detect and interpret the attack much quicker.

4.3.3.1 Simulation of the Problem

To illustrate the problems we must address, as well as a possible solution, we developed a simulated environment. We refer to this simulation throughout this section.

The simulated environment supports 100 different report types of unusual or unexpected events in our network. These may indicate a web server crashed, a connection to a non-existent server was attempted, etc. In a typical day our simulator generates 10 reports, and these are independently and randomly assigned to one of the 100 report types. The 10 reports may be assigned to 10 different report types (common) or they may all be assigned to the same type (very unlikely).

In this network an attacker is launching a very slow and subtle attack. The attack is new, so our signature-based systems do not detect it, but it does register somehow, somewhere in our sensor grid (i.e., it "jiggles" one of our wires). The attacker only launches a single instance of the attack on any given day, and he does not attack every day. Only with a probability of 50% does he try his attack on any particular day (i.e., he attacks roughly every other day).

We have also generated a very simple user interface indicating the number of reports for each event type that was observed for that day. The interface consists of a 10x10 grid (i.e., 100 squares) with each square representing one event type (e.g., "rejected connection attempt to port 109"). The more reports of a particular event type, the higher the bar graph is over of the square for that event type.

Figure 11 illustrates the interface a system administrator may see on a typical day. The reported anomalous events are randomly distributed on the report grid. In one case, two reports were generated for the same report class. The report associated with our attacker is marked in the figure; however, at this point, this particular report is indistinguishable from all the other reports.

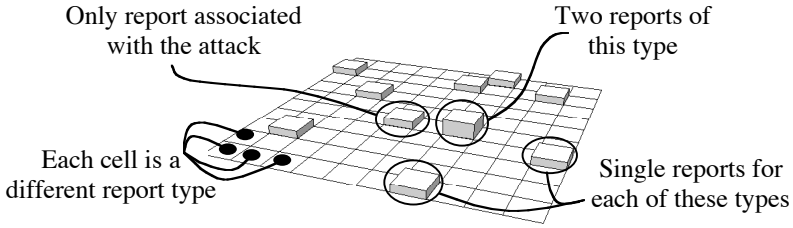

Figure 11: Interface of Unusual Events

The first problem this site will encounter is that most anomalies are benign. In any moderately sized network there are virtually an infinite number of unusual or unexpected events that can occur, and even if a very small number do occur, their numbers can easily overwhelm the reports associated with subtle attacks. At this point, the site cannot distinguish between reports associated with benign and malicious activity, so conducting an exhaustive analysis to find the underlying cause of each report must be performed.

In our simulation, on days that an attack is launched, more than 90% of reports are still caused by benign activity. On days an attack is not launched, 100% of reports are caused by benign activity.

The second problem, as mentioned previously, is that the linkage between a report and the underlying cause is often tenuous at best. Domain-specific expertise may be required (e.g., to diagnose unusual telnet protocol negotiations that may be causing a server to freeze), and in the case of transient anomalies, post-mortem analysis to determine the cause of a report may be theoretically impossible because the evidence is no longer available.

So the site is left with a number of reports, each which may require hours to diagnose, some which will be impossible to diagnose, and the underlying causes for the vast majority of these reports are benign. We should not be surprised that people tasked with tactical operations (e.g., performing the day-to-day activities to keep a site secure) tend not to collect and analyze such data.

4.3.3.2 Distinguishing Between Systematic and Transient Activity

While we are detecting the telltale evidence of the attack, we are currently unable to distinguish it from the vast majority of benign activity. However, through aggregation and statistical analysis we can greatly reduce the number of reports that must be analyzed.

Our assumption is that the attack is systematic. That is, while the attack may be new and launched as subtly as possible (i.e., “low and slow”), the attacker will eventually carry out the attack many times against the network. We exploit this attribute of the attack to our advantage by performing simple statistical analysis on large volumes of reports.

We ran our simulated network 100 days and averaged the reporting results of all the days to generate a composite graph (see Figure 12). Through this relatively straightforward analysis, we are able to clearly identify the reports associated with the attack.

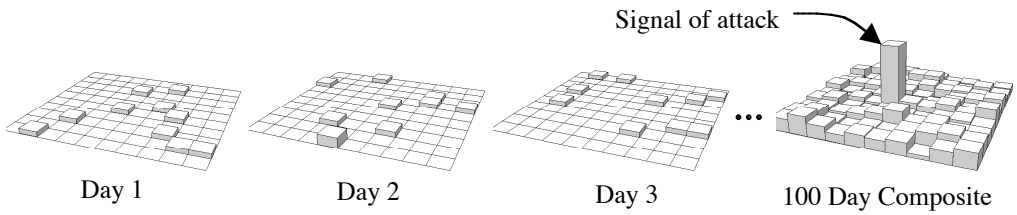

Figure 12: Aggregation of Anomalies

4.3.3.3 Interpretation and Accelerating Detection

Despite our ability to clearly identify the systematic activity through aggregation and trend analysis, one partial problem remains, and we have introduced a new problem. First, the aggregation technique does not distinguish between benign and malicious activity, rather it distinguishes between transient and systematic activity. A human must still interpret the cause behind the reports, and this may require extensive work, be beyond the expertise of a particular site, or be impossible to analyze without additional data, perhaps even including source code.

However, we have greatly reduced the amount of analysis that must be performed. In particular, during the 100 days of analysis in our simulated environment over 1000 reports were generated covering all 100 report classes, but our administrator is now focused on reports of a single class.

The new problem introduces is that 100 days were required to clearly establish the signal of systematic activity in our network. Obviously in a real-world environment the time to establish a clear signal will vary depending on the level of "background noise" a site typically generates and the patience of the attacker. The fundamental point, however, is that detection of subtle but systematic activity at a single site may require tens or hundreds of days before the activity is clearly detected.

However, new techniques are rarely used against a single site. A single attacker may apply the technique against many related sites to achieve his goal. For example, if an attacker's goal is to disrupt air operations in a specific location, the attack may be launched against multiple military bases involved in the operations, contractor sites that are supporting the operations, and even companies whose equipment is used in the operations. Also, more than one attacker, or attack group, may be using the same tool or technique for multiple and independent reasons.

A strategic intelligence group can take advantage of this behavior to accelerate detection of the signal. In particular, if the attack is launched against 100 different sites, aggregating anomaly reports across these sites can reveal the signal in a single day instead of the previous 100 days. Thus, in Figure 12 instead of each graph representing a single day at the same site, each graph would represent the same day at many different sites, and the composite graph represents the activity for one day (at 100 sites).

To summarize this section, for many practical reasons most people tasked with tactical operations do not collect and analyze the large amounts of reports that might identify a new and subtle attack. A single site can potentially detect the attack by performing long-term trend analysis on their network activity, but this approach could take days or weeks. Furthermore, since the attack is new, any individual site might not have the depth of expertise to really understand the fundamental nature of the attack responsible for the associated reports. A strategic intelligence group, on the other hand, can (1) detect a systematic attack much quicker than any individual site can, and (2) because strategic intelligence can provide a much greater depth of

expertise than any single site can, it is better suited to interpreting the nature of the attack from the reports that identified it.

5 Conclusions

This paper was motivated by discussions with a number of people over where to place network-based intrusion detection sensors. The answer depended what you wanted to do with the sensor information. This led to an examination of three types of activities that use intrusion detection sensor data: (1) tactical operations, (2) aggregated tactical operations, and (3) strategic intelligence.

Tactical operations is concerned with the hands-on activities that must be performed during the course of running a secure operation. This includes everything from installing operating systems and patches to blocking active attacks. The fundamental goal for sensor placement and configuration from a tactical operations point of view is to reduce the number of generated reports to a small handful that clearly need human attention. For this type of activity, we argue that a network-based sensor should be placed behind a firewall. We also believe there are additional steps that can greatly reduce the number of reports that need to be processed by humans. For example, combining sensor information from a well-placed sensor, vulnerability information from scanners, and data from services such as ICAT, the entire intrusion detection system can generate small numbers of actionable reports.

Aggregated tactical operations takes a subset of the tactical operations from many different sites and brings them together under a single organization. Proponents of this approach often point to economies of scale that can be achieved through such approach. Proponents of aggregation also point to the pooling of intellectual knowledge and being able to see the bigger picture than any individual site can as additional advantages. We agree to some of these benefits from aggregation, but we believe most of the true benefits fall under the mission of strategic intelligence.

Strategic intelligence is concerned with understanding the threat environment and using that knowledge to optimize tactical operations. Strategic intelligence can support predicting attacks, identifying potentially important attacks, and detecting and interpreting subtle new attacks. Strategic intelligence is not concerned with managing the typical day-to-day attacks, so reducing sensor reports to the bare minimum is not important. In fact, collecting large numbers of reports, and then processing them through statistical analysis algorithms (e.g., miscellaneous data mining algorithms) is desirable, so placing a sensor in front of the firewall may be the best choice.

6 References

[Hebe 01] Heberlein, Louis T., *Before Applying New Technologies*, TR-2001-05, Net Squared, Inc.

[Mang 99] Manganaris, S., Christensen, M., Zerkle, D., Hermiz, K., “A Data Mining Analysis of RTID Alarms,” *Proceedings of the Second International Workshop on Recent Advancements in Intrusion Detection*, Sep., 1999.

[Mukh 94] B. Mukherjee, L.T. Heberlein, K.N. Levitt., "Network Intrusion Detection," IEEE Network, Vol. 8 No. 3, pp. 26-41, May/June 1994.

[Nach 02a] Nachenberg, Carey, "Detection Avoidence", *Position Papers for the DARPA Malicious Code Defense Workshop*, Aug 2002.

[Nach 02b] Nachenberg, Carey, Private communications, Aug 2002.

[Schn 00] Schneier, Bruce. *Secrets and Lies: Digital Security in a Networked World.* New York: John Wiley & Sons, Inc., 2000.

[Schn 01] Schneier, Bruce. *Managed Security Monitoring: Network Security for the 21st Century*. Counterpane Internet Security. 2001.

[Schw 02] Schwartz, John. "Year After 9/11, Cyberspace Door Is Still Ajar," *New York Times*, 9 Sep 2002.

[WHO 99] "Principles of Disease Surveillance", World Health Organization (WHO), http://www.who.int/emc/slideshows/Survintro/sld001.htm, Oct. 1999.