

# **Review of the CPP Cyber Security Program**

Version 1.0

29 June 2005

Todd Heberlein & Tye Stallard  
Net Squared, Inc.

# Table of Contents

1	Executive Summary .....	4
2	Findings and Recommendations .....	5
2.1	Findings .....	5
2.2	Recommendations .....	6
3	CPP Cyber Security Overview .....	8
3.1	Cyber Security Monitoring in the CPP Context.....	8
3.2	DOE’s Cyber Security Monitoring And The Air Force’s .....	9
3.3	Findings and Recommendations .....	9
4	Actionable Information .....	10
4.1	Actionable Information as Organizing Principle .....	10
4.2	Simple Example: Actions Based On An Alert.....	11
4.3	Findings and Recommendations .....	12
5	Cyber Security Missions and the Role of Sensors .....	12
5.1	Introducing the Problem .....	13
5.2	Tactical Operations.....	15
5.2.1	Moving the Primary Sensor .....	16
5.2.2	Integrating Vulnerability Knowledge .....	17
5.3	Aggregated Tactical Operations.....	17
5.4	Strategic Intelligence .....	19
5.4.1	Attack Prediction .....	20
5.4.2	Identifying Important Attacks .....	22
5.4.3	Detecting New Threats .....	23
5.5	Summary of Cyber Security Missions and Use of Sensors .....	28
5.6	Findings and Recommendations .....	29
6	CPP Sensor Issues.....	30
6.1	Snort Sensor .....	30
6.2	Sensor Placement .....	30
6.3	Encryption.....	31
6.4	Findings and Recommendations .....	31
7	Conclusions.....	32
8	References .....	32

## List of Figures

Figure 1: CPP Organization and Information Flow Chart .....	8
Figure 2: Classical Sensor Placement .....	14
Figure 3: Sensor Purpose Spectrum.....	15
Figure 4: Alternative Sensor Placement.....	16
Figure 5: Integrating Vulnerability Information.....	17
Figure 6: Disease Surveillance [WHO 99].....	19
Figure 7: Myopic Sensor View.....	20
Figure 8: Over-the-Horizon Threat Detection.....	21
Figure 9: Measuring Site Similarity.....	22
Figure 10: Best Sellers vs. Uniquely Popular.....	23
Figure 11: Signature vs. Anomaly Reporting.....	24
Figure 12: Interface of Unusual Events .....	26
Figure 13: Aggregation of Anomalies .....	27

# 1 Executive Summary

This report describes our findings and recommendations for the Department of Energy's (DOE) CPP<sup>1</sup> defensive cyber security operations. The source for our information included interviews with Lawrence Livermore National Laboratory (LLNL) Computer Incident Advisory Capability (CIAC) personnel, unclassified documentation provided by CIAC, and information retrieved from the web. Our focus was on *defensive* cyber security and not on offensive operations or counterintelligence.

There are two key themes in our findings. First, the CPP sensor grid was not designed for nor is it operated in way to support defensive cyber security. CPP was initially a counterintelligence product and it seems quite obvious that the defensive cyber security mission was an after thought in its design and continues to be in its operations. Perhaps the most glaring example of this problem is the lack of any data available to a defensive cyber security analyst to (1) confirm that activity that generated an alert is really malicious and (2) determine the outcome of an attack attempt. The end result is that if a defensive cyber security analyst does try to warn a DOE site about a possible attack, there is a very strong probability either the attack never happened (false alarm) or the attack failed because the system was not vulnerable. In either case the individual site security personnel may spend hours investigating a non-event, and the defensive cyber security effort will lose credibility. There are numerous other examples of design or operations of the CPP that pose problems for a defensive cyber security mission.

The second major theme is that the defensive cyber security has not been clearly articulated. "Defensive cyber security" can encompass a wide range of missions (we describe several in Section 5), and without a specific, well-defined mission that can be reasonably accomplished within the expected budget, success or failure of the effort cannot be evaluated. For example, superficially the CPP defensive cyber security effort looks similar to the Air Force Information Warfare Center and Air Force Computer Emergency Response Team (AFIWC/AFCERT) efforts, but the Air Force efforts have a staffing level (and presumably budget) that is orders of magnitude greater than the DOE's effort. AFIWC/AFCERT has also been developing their technology and infrastructure for over a decade. Given these differences between the Air Force and DOE programs, expecting both to perform the exact same mission is unreasonable. In this report we recommend a report that the CPP defensive cyber security mission should be based on a strategic cyber intelligence capability (described in Section 5.4), and many of our findings and recommendations reflect this.

Section 2 summarizes our findings regarding the current CPP implementation and recommendations for future work. Section 3 describes the CPP sensor grid and briefly compares it to probably the most mature similar system, the Air Force's ASIM sensor grid developed by AFIWC and operated by AFCERT. Section 4 talks about actionable information and how different people within the DOE cyber infrastructure continuum will want to take different actions based on the same information because of different goals. Section 5 describes three different missions that use intrusion detection sensors and how these missions relate to choices such as sensor placement. In particular, Section 5.4 describes a cyber strategic intelligence mission, and we believe the DOE's CPP defensive cyber security could use this mission as a starting point to define their own mission. Section 6 focuses specifically on the CPP sensors. Finally, Section 7 summarizes this report.

---

<sup>1</sup> A review of DOE documents shows that CPP has at various time stood for "*Computer Protection Program*", "*Cooperative Protection Program*", and "*Cyber Protection Program*".

## 2 Findings and Recommendations

This section summarizes our findings and lists our recommendations. Each finding and recommendation comes from a specific section in the document, and that section number is listed at the end of each bullet. The same findings and recommendations are also found at the end of each section. Because a finding or recommendation may pertain to more than one section, they may be repeated (e.g., quickly deploying new signatures).

### 2.1 Findings

- **The DOE’s defensive cyber security monitoring effort is an “add on” to the CI monitoring effort.** The CPP was designed to support CI, and some of the requirements and restrictions to serve that need are inappropriate to the needs of a defensive cyber security effort. (Section 3)
- **There is no tight coupling between sensor design, signature creation and deployment, and monitoring.** These efforts have largely proceeded independently. (Section 3)
- **The CPP does not automatically collect data associated with a potential attack.** Due to restrictions from either EO 12333 or PD 61, CI personnel cannot look at data (beyond packet headers) until a case has been opened. For defensive cyber security, analysts need to look at packet data to (1) determine if an alert from a signature is a false alarm or not, (2) determine if the attack was successful, and (3) determine the underlying cause of unusual activity (is it malicious or benign). (Section 3)
- **The DOE’s defensive cyber security effort has a budget that is tiny compared to other efforts such as the Air Force’s AFIWC/AFCERT effort.** (Section 3)
- **Most DOE sites support local tactical intrusion detection capabilities.** In some cases there may be various ways to improve their operations (sensor placement, integration with vulnerability scanners), but there is already an existing capability in most places. (Section 5)
- **While the DOE does not provide a full aggregated tactical intrusion detection capability such as AFIWC/AFCERT or Counterpane, it does support an “on demand” technical support capability through CIAC.** Supporting a full Managed Security Services or Managed Security Monitoring would require a huge staffing increase. (Section 5)
- **The CPP defense cyber security could fulfill the Cyber Strategic Intelligence mission.** (Section 5)
- **Infrequent signature updates to the CPP sensors limit its ability to warn DOE sites about spreading threats before they hit most DOE systems.** If one of the goals of the CPP defensive cyber security is to warn sites, steps to make sure changes in the threat situation is detected in time to provide a timely warning to sites. (Section 5)
- **Lack of supporting anomaly detection capability in the CPP sensors limit their ability to detect new and subtle attacks.** If one of the goals of the CPP defensive cyber security is to detect new attacks, the CPP sensors should be designed to support the detection of new attacks. (Section 5)
- **Lack of “drill down” capability prevents analysts from determining the cause of anomalies.** Even if the CPP sensors supported anomaly detection, without being able to review additional data (e.g., the data associated with the anomaly), analysts would not be

able to determine if a detected anomaly is caused by malicious or benign activity. (Section 5)

- **The CPP defensive cyber security system cannot determine what attacks are unique to the DOE.** In order for the DOE to determine what threats are unique to the DOE (and therefore may pose the greatest danger), the DOE must compare the activity it is seeing with non-DOE sites. (Section 5)
- **CPP Snort sensors tend to have old signatures reducing their effectiveness at early detection.** Symantec reports that the average time between vulnerability announcement and the first exploit is less than a week [Syma 05], so CPP signatures may not be installed until well after most systems have been attacked at least once. (Section 6)
- **CPP Snort sensors apparently do not have a “tuned” signature rule set.** Snort experts rarely use the default signature snort set, which is essentially a large collection of rules submitted by a number of contributors with varying levels of skills. The large number of rules creates a burden on the sensor hardware, and the wide range of quality in the rules causes a large number of false positive. (Section 6)
- **The locations of CPP sensors are not well documented.** This may result in inappropriate filtering of event reports. Lack of knowledge of sensor placement can lead to a false picture of the range and intensity of threats a particular organization faces. (Section 6)
- **The CPP sensor grid does not appear to address encrypted activity.** Many of the DOE’s most important services (from interactive login to password-restricted web sites) are often protected by encryption that the CPP sensors cannot analyze very well. The result is that the most important cyber services are the least analyzed. (Section 6)

## 2.2 Recommendations

- *Recommendation:* **Determine a clearly articulated mission for DOE’s defensive cyber security monitoring that fits within its expected budget and for which it can bring unique qualities not available at individual sites.** We have described one such mission in Section 5.4. (Section 3)
- *Recommendation:* **Develop tighter couplings between the sensor grid design and development, signature operations, and security monitoring.** For example, a every false alarm from a signature rule should be sent back to the signature team so that they can refine their signatures. (Section 3)
- *Recommendation:* **Collect data associated with suspicious activity.** Without some level of ground truth (e.g., the data responsible for generating an alert), the defensive cyber security team cannot have any confidence in their analysis nor any guidance in how to improve their analysis. (Section 3)
- *Recommendation:* **Develop a list of roles within DOE that have security responsibility, determine what actions they need to take, and determine what specific information and under what context do they need to see it in order to trigger the actions.** (Section 4)
- *Recommendation:* **Identify what actionable information the DOE defensive cyber security can deliver and to whom.** In other words, who are the customers and how does CIAC/CS maximize the value they provide to them. (Section 4)

- *Recommendation: **Determine where actions may conflict and develop conflict resolution guidelines.*** In some cases, different players' goals can be at odds with one another. Quickly securing a penetrated machine can interfere with CI's efforts to determine who penetrated the system. (Section 4)
- *Recommendation: **Update signatures frequently.*** The value of a signature is optimal when there are a number of vulnerable systems in the network. Once all systems are either patched or successfully attacked, detecting additional attacks against those systems provides only marginal value. (Section 5)
- *Recommendation: **Include additional technology in CPP sensors to support anomaly detection.*** This could include additional summary of content (e.g., checksum on data for first several packets) to sensor-side histories (e.g., the DOE Network Intruder Detector included session path anomaly score in the session data). (Section 5)
- *Recommendation: **Develop "session signatures" of systems that were successfully attacked.*** Sometime once an attacker has initially penetrated a host, he connects from the penetrated hosts back to some system he controls to download additional tools such as rootkits. Another common behavior is for the attacker to send a message to a chatroom (e.g., a specific IRC channel) to brag about his success. Once the penetration of a DOE system has been detected through whatever means, these outbound connection patterns should be identified (the "session signatures"). Then CPP defensive cyber security analysts should search their database for past evidence of such "session signatures" and keep an eye out for future evidence of such activity. This is a relatively unique capability that the CPP data offers that many commercial systems do not. (Section 5)
- *Recommendation: **Record additional information to allow CPP defensive cyber security analysts to "drill down" into the data.*** The "drill down" capability is necessary to determine whether an anomaly represents a new threat that the DOE should be concerned with. (Section 5)
- *Recommendation: **Exchange threat profiles with other government agencies to determine threats that are unique to the DOE.*** This activity includes (1) summarizing and sanitizing the data so that the other government agencies cannot determine which DOE sites have been attacked; (2) developing standard data structures (perhaps in XML) to represent the data; (3) developing network protocols to support automatic exchange of the data; and (4) establishing Memorandum Of Understanding between participating organizations. While exchanging information is politically sensitive, being able to identify attacks that are unique to the DOE may help identify the most important attacks. (Section 5)
- *Recommendation: **Create a signature team for the CPP.*** The team would provide quality assurance testing and refinement on signatures, identify the most appropriate signatures to deploy in the CPP sensors, and develop signature for newly detected attacks. (Section 6)
- *Recommendation: **Document the DOE site's network architecture and the sensor placements in this architecture.*** (Section 6)
- *Recommendation: **Create a taskforce to determine how to analyze attacks that current CPP sensors cannot analyze.*** The most obvious cases involve encrypted services, but NAT, internal switched networks, and other trends pose challenges to the current generation of sensors. (Section 6)

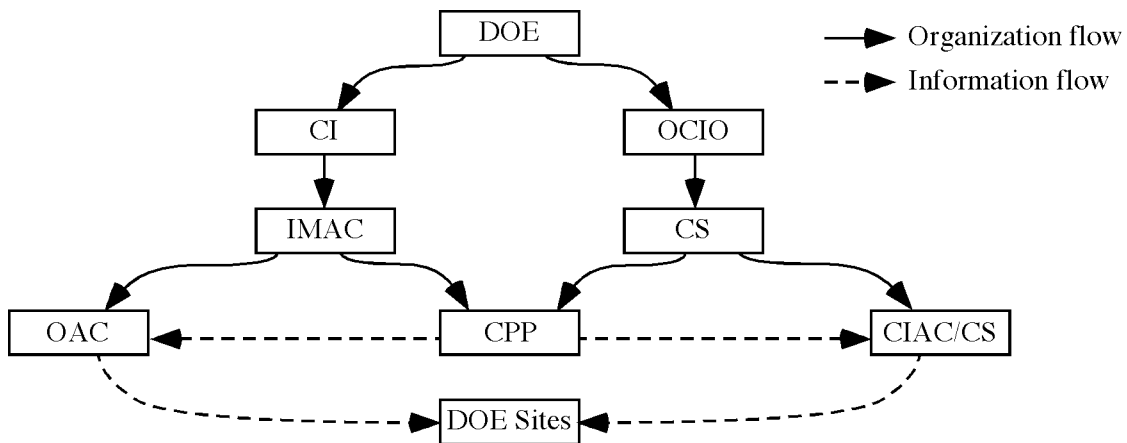
### 3 CPP Cyber Security Overview

This section summarizes the CPP effort and its relation to the Cyber Security effort. It provides the context for how we evaluated the effort. We begin by looking at the historical and organizational context for the CPP cyber security monitoring effort, and we follow that by comparing the DOE’s effort with the Air Forces, probably the oldest and most mature monitoring effort.

#### 3.1 Cyber Security Monitoring in the CPP Context

Figure 1 shows the basic organizational and notional information flow for the Cooperative Protection Program (CPP). At the top is the Department of Energy (DOE), under which there are the offices of Counterintelligence (CI) and Office of the Chief Information Officer (OCIO). The Office of the CI supports the Inquiry Management and Analysis Capability (IMAC), which supports the CPP and the Operational Analysis Center (OAC) at the Pacific Northwest National Laboratory (PNL). The OCIO supports Cyber Security (CS), which supports CPP and Computer Incident Advisory Capability Cyber Security effort (CIAC/CS) at Lawrence Livermore National Laboratory (LLNL).

Notionally the CPP represents a DOE-wide cyber sensor grid which provided data to the Counterintelligence’s OAC and Cyber Security’s CIAC/CS. These organizations provide various levels of data and services to the individual DOE sites. For example, CIAC/CS runs the “AWARE” web site, and individuals at DOE sites can retrieve summaries of CPP information through this web site. In practice, the CPP data actually goes to the IMAC/OAC site at PNL and then to CIAC/CS at LLNL.



**Figure 1: CPP Organization and Information Flow Chart**

The flow of the data from the CPP sensor grid to PNL and then LLNL is part historical and part political. The original CPP sensor grid was developed for and sponsored by CI, so initially the right side of Figure 1 did not exist. The CPP, started by Presidential decision directive (PDD) 61 [DOE 00], was designed and developed to support the Counterintelligence efforts defined by Executive Order 12333 [EO 12333] as it pertains to cyberspace. DOE sites are required to participate in CPP for CI purposes (although not for Cyber Security purposes). EO 12333 restricts the ability of CI personnel from looking at data until a case is actually opened, because of this restriction the CPP sensor grid was designed *not* to collect any packet data beyond the headers. This lack of data has profound implications on CIAC/CS ability to perform their cyber security duties.



Later, the OCIO's Cyber Security effort also provided money to support the CPP sensor development and deployment as well as a CIAC/CS analysis capability of the CPP data. Whereas DOE sites are required to participate in the CI effort, they are not required to participate in the Cyber Security effort, thus OAC receives sensor feeds from more sites than the CIAC/CS does.

In an extreme oversimplification of the difference between CI and CS missions, CI can be seen as determining who is doing what to whom and CS can be seen as simply hardening the cyber infrastructure against malicious acts irregardless of who is performing it. While the OAC does perform regular analyses of the data, its primary role does not begin until a case is opened. CS, on the other hand, has a continuous role in helping the DOE harden their systems on a daily basis. These different goals can lead to different responses to observed activities, and in some cases these responses can potentially conflict with each other. Section **Error! Reference source not found.** describes some of the different responses that can be taken in response to an alert from a CPP sensor.

### 3.2 DOE's Cyber Security Monitoring And The Air Force's

The CIAC/CS portion of the CPP effort is superficially similar to the Air Force's ASIM effort. The Air Force Information Warfare Center (AFIWC) develops and supports sensors and analysis tools for the Automated Security Incident Measurement and Common Intrusion Detection Director System (ASIM/CIDDS). The Air Force Computer Emergency Response Team (AFCERT) monitors the Air Force wide ASIM sensor grid 24x7.

However, the AFIWC/AFCERT differs from the CIAC/CS effort in a couple of critical ways. First, for the Air Force, AFIWC/AFCERT creates an extremely tight coupling between sensor development and deployment, signature creation, and monitoring – it is all essentially done under one roof (figuratively speaking). This tight coupling creates fast and strong feedback loops for continuous improvement. For the DOE one group does sensor development and deployment (e.g., PNL), one group does signature generation (the open source community), and one group does the monitoring (CIAC). The feedback loops are weak or non-existent.

A second difference is that the DOE has a much wider ranging user community with different expectations than the Air Force has. Some DOE sites are tightly run like a military base, but many sites more closely resemble a university environment. This difference produces a much wider range in network behavior than the military has, and this variability makes the DOE monitoring effort much more difficult.

A third difference is that AFIWC/AFCERT has a huge operational budget compared to the DOE's CIAC monitoring budget. For example, AFCERT has maintained a staff of several dozen full time analysts. Historically AFIWC/AFCERT has supported an aggregated tactical intrusion detection operation (see Section 5.3). Without a tremendous budget increase, the DOE cannot expect to match the Air Force capabilities, so an alternative mission should be considered. Section 5.4 explores one mission, strategic cyber intelligence, that may make for a strong fit the DOE's cyber security monitoring effort.

### 3.3 Findings and Recommendations

- **The DOE's defensive cyber security monitoring effort is an "add on" to the CI monitoring effort.** The CPP was designed to support CI, and some of the requirements and restrictions to serve that need are inappropriate to the needs of a defensive cyber security effort.
- **There is no tight coupling between sensor design, signature creation and deployment, and monitoring.** These efforts have largely proceeded independently.

- **The CPP does not automatically collect data associated with a potential attack.** Due to restrictions from either EO 12333 or PD 61, CI personnel cannot look at data (beyond packet headers) until after a case has been opened. For defensive cyber security, analysts need to look at packet data to (1) determine if an alert from a signature is a false alarm or not, (2) determine if the attack was successful, and (3) determine the underlying cause of unusual activity (is it malicious or benign).
- **The DOE’s defensive cyber security effort has a budget that is tiny compared to other efforts such as the Air Force’s AFIWC/AFCERT effort.**
  - *Recommendation: Determine a clearly articulated mission for DOE’s defensive cyber security monitoring that fits within its expected budget and for which it can bring unique qualities not available at individual sites.* We have described one such mission in Section 5.4.
  - *Recommendation: Develop tighter couplings between the sensor grid design and development, signature operations, and security monitoring.* For example, a every false alarm from a signature rule should be sent back to the signature team so that they can refine their signatures.
  - *Recommendation: Collect data associated with suspicious activity.* Without some level of ground truth (e.g., the data responsible for generating an alert), the defensive cyber security team cannot have any confidence in their analysis nor any guidance in how to improve their analysis.

## 4 Actionable Information

One of the drawbacks of today’s information age is information overload. People are regularly presented with information that they do not need, do not need in that form, or do not need right at that moment. In the field of intrusion detection this is a common problem. Users are often presented with volumes of information that they find useless and must sift the wheat from the chaff. This reduces productivity, creates burnout, and tends to create high turnover rates at security operations centers. For example, we often hear that AFCERT has a very high turnover rate. The DOE security efforts are no different, so this section briefly examines these issues.

### 4.1 Actionable Information as Organizing Principle

To the greatest extent possible, information presented to a human should be actionable. For example, if a sensor detected a known attack against a known vulnerability but that was blocked by a firewall, this alert in and of itself should not be presented to a human because there nothing the human needs to do. However, if there is a known attack against a system known to be vulnerable, this alert should be presented to a human because human action needs to be taken.

However, whether a piece of information is actionable depends on who receives that information, what his role within the organization is, and what the goals of the organization are. For example, if an analyst is a Quality Assurance engineer for signatures, he may initially want to view all alerts triggered by a rule in order to verify that the rule fired on proper data; after a while he may only want to sample data; and eventually he may not want to see any alerts from that rule. Meanwhile, a network or system administrator may only want to see the alert when it is associated with a *vulnerable* system under his responsibility (i.e., the attack is probably successful), and he will always want to see such alerts.

Some people will want to see different information. Some people will want to see the same information at different times (depending on additional context). People seeing the same

information may want to take different actions depending on what their goals are. Determining who the various players in the DOE security continuum are (users, system administrators, signature engineers, CI, CIAC strategic analysts, etc.), what types of actionable information they want to see, and what actions they would take is important to understanding how a DOE defensive cyber security global monitoring capability fits into the bigger scheme. And understanding this will guide the development of the DOE defensive cyber security monitoring effort.

## 4.2 Simple Example: Actions Based On An Alert

In this section we examine some of the types of actions that can occur once a sensor generates an alert. For this discussion we use the term “analyst” to anyone who may be tasked with taking some action following the alert: this may include global analyst looking for strategic trends, local network analysts, system and network administrators, and even users. The key is that these are actions to satisfy some goal that take place only *after* an alert has been generated:

- **Verify there was an attack and determine if it was successful.** Unfortunately intrusion detection systems tend to generate large numbers of false alarms, so an alert generated by a sensor often needs to be investigated for accuracy. Furthermore, assuming an organization updates their operating systems, applications, and anti-virus software on a relatively regular basis, most observed attacks should fail. Unless the sensor system is tightly integrated into vulnerability information, an analyst must invest time determining if an attack was successful or not. A successful attack requires additional work. An alert from a sensor must be verified, and someone must determine if further analysis is required.
- **Broadcast an alert about the attack.** If the attack was a new attack or against a new vulnerability, or if the attack is part of an escalating trend, analysts may want to notify other organizations about the attack so they can harden their systems before being attacked.
- **Return system and users to operation as soon as possible.** Once an organization detects that malicious code (e.g., a worm or virus) or an interactive adversary (e.g., a hacker) has penetrated a system, the system is typically removed from operation for some period of time. Unfortunately this means that the computer system is not being used to support the organization. If it is part of an e-commerce site, the system is not generating revenue. If the system is part of a research project, the research project is slowed down. If the system is supporting an on-going operational activity such as supporting a war effort, the operational activity is reduced. If the system was an employee’s workstation, the employee is often idle. Because of these reasons organizations often press for a system to be cleaned and restored to operational capability as soon as possible.
- **Determine what information was compromised or corrupted.** Depending on the nature of the penetrated information system, an analysis must be performed to determine what information has been compromised or what information was corrupted. For example, if the computer system contained organizationally sensitive information (e.g., Classified documents), identifying what information was (or may have been) compromised must be determined, and additional actions may need to be taken (e.g., generating reports, notifying users that passwords were compromised, etc). In addition to compromised information, some data may be compromised – applications may be Trojaned, database information may be modified, words in documents could be changed.
- **Determine how the attack occurred.** To prevent repeated penetrations into the system once it is restored to service, and to prevent penetrations of other machines with the same vulnerability, analysts must determine how the adversary penetrated the system and how

the system can be hardened (e.g., applying a patch). If the penetration was by a known attack for which there is a known patch, organizational processes may need to be analyzed to determine the decision process that allowed the vulnerability to remain open. If the penetration was by an unknown attack against an unknown vulnerability, a technical investigation may need to be launched that can require specialized knowledge of the various technologies (operating system, application details, programming language issues, network protocols, etc.).

- **Collect information for prosecution.** Many organizations may wish to prosecute the person responsible for the penetration. In such situations, evidence must be collected and preserved according to appropriate rules of evidence. Frequently this involves law enforcement agencies confiscating the hard disk or the entire computer system.
- **Collect information for counter intelligence.** In order to capture the individual responsible for the penetration, or to determine if the person was after some specific information, the appropriate analysts may want to leave the system “in place”, in order to collect additional information about the adversary. This can include tracing back the adversary to his point of origin or simply observing what types of information the adversary is searching.

Unfortunately, many of these actions can interfere with one another. For example, quickly returning a system to full operational status may involve scrubbing the system, reinstalling the operating system, applications, and data, and applying all appropriate patches. These actions, especially if taken quickly, can interfere with properly collecting and preserving data for prosecution. Similarly, if a new attack is detected and a broadcast is widely distributed to the community, the adversary may realize he has been detected and change his tactics, thus interfering with counter intelligence efforts.

Furthermore, since these actions are often associated with a wide range of people in widely distributed locations, determining who should take which actions and when (or under what circumstances) presumes a clear understanding of lines of authority and responsibility, established rules engagement, and established lines of communications.

### 4.3 Findings and Recommendations

- *Recommendation:* **Develop a list of roles within DOE that have security responsibility, determine what actions they need to take, and determine what specific information and under what context do they need to see it in order to trigger the actions.**
- *Recommendation:* **Identify what actionable information the DOE defensive cyber security can deliver and to whom.** In other words, who are the customers and how does CIAC/CS maximize the value they provide to them.
- *Recommendation:* **Determine where actions may conflict and develop conflict resolution guidelines.** In some cases, different players’ goals can be at odds with one another. Quickly securing a penetrated machine can interfere with CI’s efforts to determine who penetrated the system.

## 5 Cyber Security Missions and the Role of Sensors

In this section we look at different cyber security defensive missions (offensive operations are not considered) and some of the sensor requirements for those missions. Our goal is to illustrate some of the goals the CIAC/CS effort may want to consider (we recommend

Strategic Cyber Intelligence, Section 5.4). In particular, we consider three major defensive security missions:

- **Tactical Intrusion Detection.** 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.
- **Aggregated Tactical Intrusion Detection.** Aggregated tactical operations bring many of the activities of local tactical intrusion detection mentioned above 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.
- **Strategic Cyber 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)<sup>2</sup> and National Center for Infectious Diseases (NCID)<sup>3</sup> as well as the World Health Organization's (WHO) Communicable Disease Surveillance and Response (CSR) group<sup>4</sup>.

## 5.1 Introducing the Problem

Figure 2 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-time 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. DARPA program managers frequently cite similar stories of overwhelming data flows from intrusion detection sensors as a reason why DARPA must develop new technologies. CIAC/CS problems can be even worse since the CPP contains both alerts and session data. Currently CIAC/CS receives approximately 6 billion event records per day [Schr 05].

---

<sup>2</sup> <http://www.cdc.gov/epo/>

<sup>3</sup> <http://www.cdc.gov/ncidod/>

<sup>4</sup> <http://www.who.int/emc/>



**Figure 2: 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.

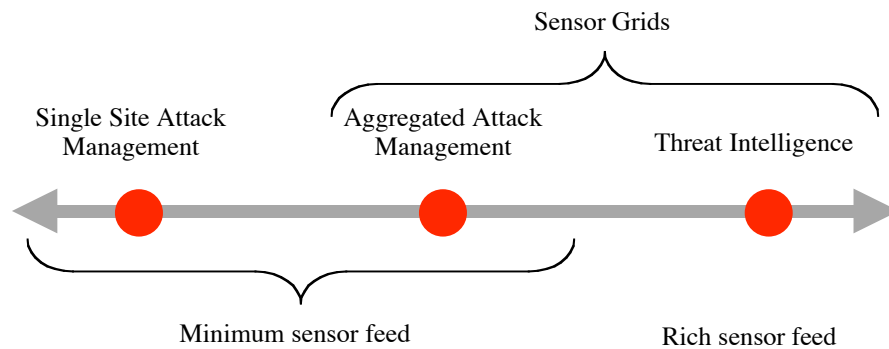
Some intrusion detection management systems provide statistical roll-up capabilities to group many alerts into a single report for display purposes; however, this does not necessarily solve the fundamental problem. If host *bad\_guy* launches an RPC attack *xdr\_overflow* against 16,000 hosts within a site, an intrusion detection sensor 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.

The report “Before Applying New Technologies” [Hebe 01] argues 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 5.2.1 and 5.2.2 present two of the suggestions from that report.

But over the last decade that the role of intrusion detection has been changing. Today there is a greater emphasis on creating large-scale sensor grids such as CPP, 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 3 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 3: Sensor Purpose Spectrum**

In this report we look at each of these roles. In Section 5.2, Tactical Operations, we look at sensors used in single site attack management. If we use war as an 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. Many (probably most) DOE sites provide their tactical IDS capability.

In Section 5.3, 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. This is often called “managed security services”, and to some degree CIAC has provides a limited version of this through their trouble ticket incident support capability.

In Section 5.4, 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.

## 5.2 Tactical Operations

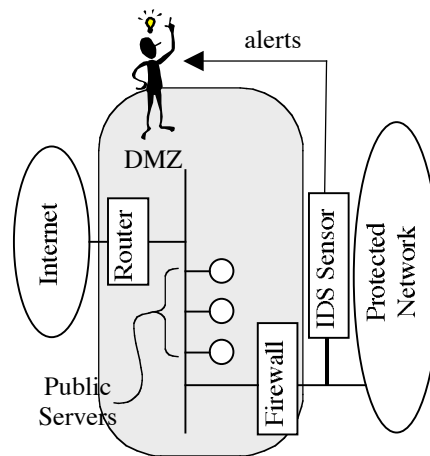
Tactical operations is concerned with managing individual attacks against a site as well as preparing a site prior to an attack. It essentially encompasses all actions by security, network, and system administrators that change the state of the network with respect to security. This can include installing 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.

### 5.2.1 Moving the Primary Sensor

The simplest and probably the most effective step to reducing the volumes of reports analysts must process is by placing the sensor behind the firewall (see Figure 4). 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 4: Alternative Sensor Placement**

There are at least three common reasons why organizations continue to place IDS sensors in front of a site's 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 the 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 the 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 organization is monitoring the site. 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 5.4.

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 be handled by tactical response teams.

## 5.2.2 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 5). 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 5: 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.

## 5.3 Aggregated Tactical Operations

Aggregated tactical operations bring many of the activities mentioned in Section 5.2 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, several ASIM architects 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 5.4.

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 Century" 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.

On the other hand, many smaller sites have system or network administrators with limited security skills, so outsourcing this capability can be important. Also, even sites with skilled system administrators might wish to have a skilled full time team provide support to reduce workload, help during reduced labor availability (e.g., when some administrators are on vacation), or provide support for analyzing attacks that take advantage of a technical areas for which the system administrators are not particularly skilled. CIAC provides such a service to members of the DOE community – if a local system administrator finds an attack for which he wants CIAC assistance he only needs to place "CIAC URGENT" in the subject line of an email to CIAC or call their hotline (925-422-8193). In a sense, CIAC provides an "on demand" tactical aggregated intrusion detection capability.

To summarize this section, while a centralized security organization has value (see Section 5.4), 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. However, in a number of cases (e.g., monitoring small sites without security expertise or requesting additional technical support), aggregated intrusion detection can be invaluable.

## 5.4 Strategic Intelligence

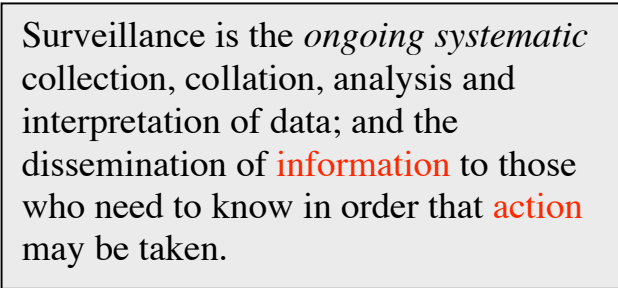
Strategic intelligence is concerned with understanding the threat environment in order to optimize tactical operations. This role is similar to the mission of the United States Center for Disease Control and Prevention's (CDC) Epidemiology Program Office (EPI)<sup>5</sup> and National Center for Infectious Diseases (NCID)<sup>6</sup> as well as the World Health Organization's (WHO) Communicable Disease Surveillance and Response (CSR) group<sup>7</sup>.

Whereas tactical operations focus 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 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 affecting 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 (e.g., by applying patches), 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 6.



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 6: Disease Surveillance [WHO 99]**

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

<sup>5</sup> <http://www.cdc.gov/epo/>

<sup>6</sup> <http://www.cdc.gov/ncidod/>

<sup>7</sup> <http://www.who.int/emc/>

they may face, so part of the role of a strategic intelligence group is to provide backup experts to help these sites. Sometimes this expertise is kept in-house at the monitoring organization, but often the expertise is employed elsewhere and is essential “on call” when needed.

Sections 5.4.1 through 5.4.3 examine some of the unique capabilities that a strategic intelligence organization can provide.


### 5.4.1 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]. In the second half of 2004, Symantec tracked an average of 54 new vulnerabilities each week [Syma 05]. 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 5.2.2).


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 (nor is it desirable since patches occasionally introduce their own problems), 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 7 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 7: 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 8). 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 8: 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.

### 5.4.1.1 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<sup>8</sup> and Amazon’s “Top 100 Bestsellers” list<sup>9</sup>.

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<sup>10</sup> and Amazon’s “Movers & Shakers”<sup>11</sup>.

A critical requirement for this to work is to have very early detection of an emerging threat so that most sites can be warned before they are attacked. Warning a site after it has been attacked does the site little good. Symantec reports that when a vulnerability is exploited, the average time between vulnerability announcement and first exploit is less than one week [Syma 05]. The CPP currently updates its deployed signatures very infrequently (on the order of months between updates), so most systems have probably been targeted by the attack before the CPP has the ability to detect the attack. For the CPP to provide a viable “warn” capability, this must be addressed.

<sup>8</sup> <http://isc.incidents.org/top10.html>

<sup>9</sup> <http://www.amazon.com/exec/obidos/tg/browse/-/549066/hot/1/103-7992537-9327010>

<sup>10</sup> <http://isc.incidents.org/trends.html>

<sup>11</sup> [http://www.amazon.com/exec/obidos/tg/new-for-you/movers-and-shakers/-/books/ref=pd\\_gw\\_msgr/103-7992537-9327010](http://www.amazon.com/exec/obidos/tg/new-for-you/movers-and-shakers/-/books/ref=pd_gw_msgr/103-7992537-9327010)

### 5.4.1.2 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 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 9 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.

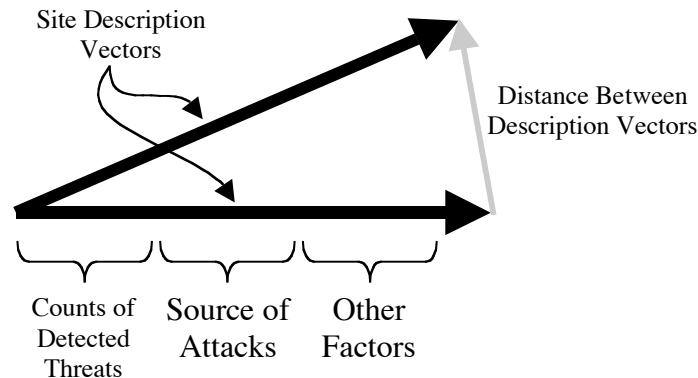


Figure 9: Measuring Site Similarity

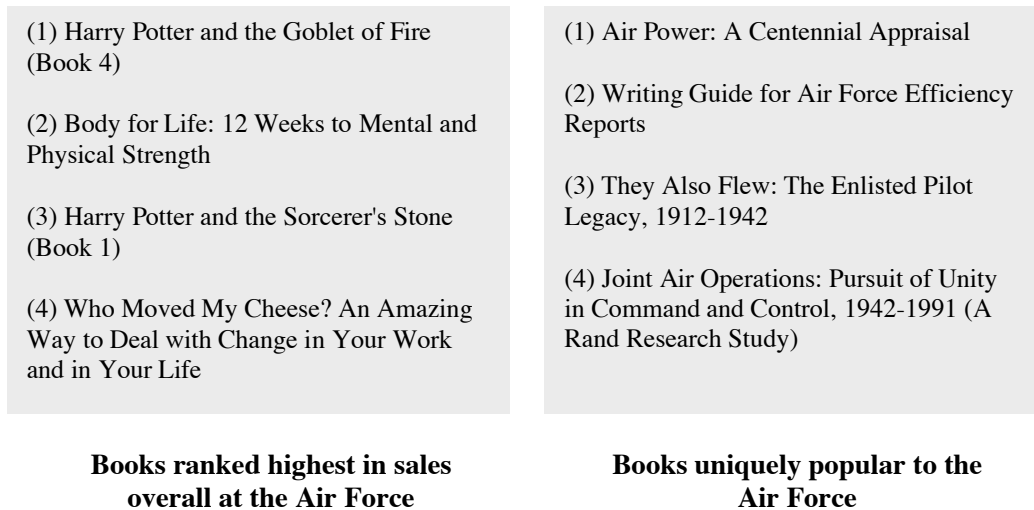
### 5.4.2 Identifying Important Attacks

There is a difference between the most active 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"<sup>12</sup>. Figure 10 illustrates this with book purchases from several years ago. 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 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.



**Figure 10: Best Sellers vs. Uniquely Popular**

In order for the DOE to determine what attacks are relatively unique to the DOE, it needs to compare its observed activity to a wider body of attack information. Beyond technical challenges, the political obstacles for such an endeavor can be enormous. For example, while DOE sites are required to send their CPP data to CI analysis, they are not required to send their data to CIAC for defensive cyber security analysis, and many do not. And this is just within the DOE. To share that information with those outside the DOE will be even more difficult.

### 5.4.3 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.

<sup>12</sup> [http://www.amazon.com/exec/obidos/subst/community/community.html/ref=gw\\_hp\\_ls\\_1\\_10/103-7992537-9327010](http://www.amazon.com/exec/obidos/subst/community/community.html/ref=gw_hp_ls_1_10/103-7992537-9327010)

We believe the primary reason for signature-based systems domination of the intrusion detection market is that in most tactical operations (see Section 5.2), 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.

Furthermore, if the DOE’s CPP defensive cyber security goals include early warning of emerging attacks, including previously unknown attacks, before most sites encounter the attacks, anomaly detection technology should play an important role. Unfortunately the CPP sensor suite does not lend itself to anomaly detection and diagnosis. Should the CPP cyber defensive security efforts consider “warn” as an operational goal, it should revisit the sensor suite to support this capability.

In Section 5.4.3.1 we look into the problem of anomaly detection for tactical operations. In Section 5.4.3.2 we briefly discuss a limitation in antivirus software and how strategic intelligence operations may address this problem. In Section 5.4.3.3 we look into detail of how anomaly detection can be used to identify subtle new attacks.

### 5.4.3.1 Problem with Anomaly Detection

As mentioned in Section 5.2, 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<sup>13</sup>) 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<sup>14</sup> 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 11, column A.

<b>Target:</b> 128.131.7.2 : 161	<b>Target:</b> 128.131.7.2 : 161
<b>Attacker:</b> 128.120.56.31 : 5611	<b>Attacker:</b> 128.120.56.31 : 5611
<b>Attack Name:</b> xdr_router_crash	<b>Attack Name:</b> unknown
<b>Vulnerability ID:</b> CVE-2002-0391	<b>Vulnerability ID:</b> unknown
<b>Vulnerable:</b> Yes	<b>Vulnerable:</b> unknown
<b>Damage:</b> Crashes Cisco routers	<b>Damage:</b> unknown
<b>Link to Patch:</b> <a href="#">Cisco_patch</a>	<b>Link to Patch:</b> none
<b>Details:</b> <a href="#">Security Focus</a> <a href="#">CERT CC</a>	<b>Details:</b> none
<b>A</b>	<b>B</b>

**Figure 11: Signature vs. Anomaly Reporting**

<sup>13</sup> <http://cve.mitre.org/>

<sup>14</sup> <http://icat.nist.gov/icat.cfm>



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 11, 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.

#### **5.4.3.2 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 "*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 DOE 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 a CIAC strategic intelligence unit, where the analysts are cleared for highly sensitive data (e.g., they should be cleared to analyze all "maybe" infected files that transit the email system). 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 CIAC 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 5.4.3.3. 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*.

#### **5.4.3.3 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.

##### **5.4.3.3.1 Simulation of the Problem**


To illustrate the problems that must be addressed, 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 signature-based systems do not detect it, but it does register somehow, somewhere in the sensor grid (i.e., it “jiggles” one of the 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 12 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 12: 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 the 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.

#### 5.4.3.3.2 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 the simulated network 100 days and averaged the reporting results of all the days to generate a composite graph (see Figure 13). Through this relatively straightforward analysis, we are able to clearly identify the reports associated with the attack.




Figure 13: Aggregation of Anomalies

#### 5.4.3.3.3 Interpreting 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 the simulated environment over 1000 reports were generated covering all 100 report classes, but an analyst is now focused on reports of a single class.

The new problem introduced 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 reveal information about a specific DOE project, the attack may be launched against

multiple DOE sites involved in the projects, contractor sites that are supporting the project, and even companies whose equipment is used for the project. 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 13 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.

## **5.5 Summary of Cyber Security Missions and Use of Sensors**

Section 5 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. This approach is also useful for supporting smaller sites without system and network administrators with security skills. A variation on this approach could be called “Aggregated Tactical Operations Lite”, where the central organization is not directly responsible for daily security operations, but they can be called upon to lend additional support once an event is detected at the local site. CIAC provides a capability similar to this within the DOE.

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. The mission of cyber strategic intelligence is probably the optimal strategy for the DOE's CPP defensive cyber security mission.

## 5.6 Findings and Recommendations

- **Most DOE sites support local tactical intrusion detection capabilities.** In some cases there may be various ways to improve their operations (sensor placement, integration with vulnerability scanners), but there is already an existing capability in most places.
- **While the DOE does not provide a full aggregated tactical intrusion detection capability such as AFIWC/AFCERT or Counterpane, it does support an “on demand” technical support capability through CIAC.** Supporting a full Managed Security Services or Managed Security Monitoring would require a huge staffing increase.
- **The CPP defense cyber security could fulfill the Cyber Strategic Intelligence mission.**
- **Infrequent signature updates to the CPP sensors limit its ability to warn DOE sites about spreading threats *before* they hit most DOE systems.** If one of the goals of the CPP defensive cyber security is to warn sites, steps to make sure changes in the threat situation is detected in time to provide a timely warning to sites.
- **Lack of supporting anomaly detection capability in the CPP sensors limit their ability to detect new and subtle attacks.** If one of the goals of the CPP defensive cyber security is to detect new attacks, the CPP sensors should be designed to support the detection of new attacks.
- **Lack of “drill down” capability prevents analysts from determining the cause of anomalies.** Even if the CPP sensors supported anomaly detection, without being able to review additional data (e.g., the data associated with the anomaly), analysts would not be able to determine if a detected anomaly is caused by malicious or benign activity.
- **The CPP defensive cyber security system cannot determine what attacks are unique to the DOE.** In order for the DOE to determine what threats are unique to the DOE (and therefore may pose the greatest danger), the DOE must compare the activity it is seeing with non-DOE sites.
  - *Recommendation: Update signatures frequently.* The value of a signature is optimal when there are a number of vulnerable systems in the network. Once all systems are either patched or successfully attacked, detecting additional attacks against those systems provides only marginal value.
  - *Recommendation: Include additional technology in CPP sensors to support anomaly detection.* This could include additional summary of content (e.g., checksum on data for first several packets) to sensor-side histories (e.g., the DOE Network Intruder Detector included session path anomaly score in the session data).
  - *Recommendation: Develop “session signatures” of systems that were successfully attacked.* Sometime once an attacker has initially penetrated a host, he connects from the penetrated hosts back to some system he controls to download additional

tools such as rootkits. Another common behavior is for the attacker to send a message to a chatroom (e.g., a specific IRC channel) to brag about his success. Once the penetration of a DOE system has been detected through whatever means, these outbound connection patterns should be identified (the “session signatures”). Then CPP defensive cyber security analysts should search their database for past evidence of such “session signatures” and keep an eye out for future evidence of such activity. This is a relatively unique capability that the CPP data offers that many commercial systems do not.

- *Recommendation:* **Record additional information to allow CPP defensive cyber security analysts to “drill down” into the data.** The “drill down” capability is necessary to determine whether an anomaly represents a new threat that the DOE should be concerned with.
- *Recommendation:* **Exchange threat profiles with other government agencies to determine threats that are unique to the DOE.** This activity includes (1) summarizing and sanitizing the data so that the other government agencies cannot determine which DOE sites have been attacked; (2) developing standard data structures (perhaps in XML) to represent the data; (3) developing network protocols to support automatic exchange of the data; and (4) establishing Memorandum Of Understanding between participating organizations. While exchanging information is politically sensitive, being able to identify attacks that are unique to the DOE may help identify the most important attacks.

## 6 CPP Sensor Issues

### 6.1 Snort Sensor

The architecture of the standard Snort software package is intended to provide a platform for experts to search network traffic for specific patterns. It is common knowledge among Snort users that the default rule set is unusable for a network connected to the Internet. It produces many false positives. Existing CPP sensors have a Snort signature set of unknown vintage. Old signatures will continue to alert on old attacks, but as application are upgraded and networks are firewalled, this information is useful only in an academic sense. In addition, the rate of dropped packets is unknown. High network bandwidth, common to DOE networks, is difficult to monitor even with the most optimized hardware based network sensor solutions.

For the purpose of forecasting Internet threats, CPP sensors need to be reconfigurable and retargetable. More than once, the Snort users community developed effective signatures for new threats via email and web site postings. Those new signatures are only useful if installed on CPP sensors while still relevant. Beyond common Internet threats, the DOE monitors for specific traffic of interest. Once a modus operandi of a specific threat is learned, wide deployment of the corresponding signature must occur. These signature sets should be issued only to authorized users of course. Attackers will attempt to avoid detection by sensors. When they can't avoid detection, they will blind them with too much irrelevant data.

### 6.2 Sensor Placement

Network sensor data has value only in relation to its network environment. Correlation of "ground truth" at the host level with general trends DOE-wide is impossible with access only to network flows. The effects on individual hosts and internal networks must be identified so that their causes may be communicated to similar hosts and networks.

The determination between attack attempts or successes can only be known if the sensor's placement on the network is known as well. A sensor outside the access control perimeter will record all known attack attempts. Whereas one inside the access control perimeter will identify those that penetrated successfully or those with an internal origin. This difference in sensor placement is significant if the analyst is monitoring an attacker over many DOE sites.

With neither a DOE standard network perimeter configuration nor a description of a sensor's environment, data sets from different sensors cannot be compared intelligently. Currently some CPP sensors monitor thousands of network addresses, others monitor only a handful (still other Internet connections into a DOE site have no sensor feeds to CIAC whatsoever (e.g., DOE HQ has several network feeds, of which only one is monitored). Also there may exist network architectures in which inbound network traffic may be on a different wire than outbound traffic for the same network session known as asymmetric routing. Just like listening to one side of a telephone conversation, detecting anything other than the most simplistic attack attempts is currently not possible.

Network scans can generate large numbers of network sessions, and CIAC was receiving approximately 6 billion session events per day from the CPP sensors. To reduce this volume to a more manageable number, CIAC chose to filter out any session with traffic that only flowed in one direction on the assumption that these represent failed connection attempts. This reduced their daily database updates from 6 billion entries per day to 2 billion entries per day, a number that significantly increased CIAC's ability to store and analyze several weeks of data. Unfortunately, if a DOE site's network is designed so that inbound traffic flowed on one wire and outbound traffic flows on a different wire – information that is not available to CIAC – then CIAC's "failed session" assumption is incorrect and potentially information is discarded.

### 6.3 Encryption

Finally, CPP data feeds are being undermined by another cybersecurity measure: encryption. This is a limitation of any network traffic analyzer. Encrypted sessions between an application server and its client, such as SSL utilized by HTTPS or SSH for remote login, prevent any examination of attempted or successful attacks on network services. Virtual private networks (VPNs) and link level encryption have the same effect over a wider range of network addresses with the additional side effect of obfuscating traffic analysis. Cybersecurity analysts cannot identify what data was transferred and cannot tell which hosts did so either.

Network address translation has a similar effect. Whereas the network packet content is not scrambled, the addresses are translated. Usually done for the purpose of efficient address space utilization, this common network management technique obfuscates what hosts transferred data.

### 6.4 Findings and Recommendations

- **CPP Snort sensors tend to have old signatures reducing their effectiveness at early detection.** Symantec reports that the average time between vulnerability announcement and the first exploit is less than a week [Syma 05], so CPP signatures may not be installed until well after most systems have been attacked at least once.
- **CPP Snort sensors apparently do not have a "tuned" signature rule set.** Snort experts rarely use the default signature snort set, which is essentially a large collection of rules submitted by a number of contributors with varying levels of skills. The large number of rules creates a burden on the sensor hardware, and the wide range of quality in the rules causes a large number of false positive.

- **The locations of CPP sensors are not well documented.** This may result in inappropriate filtering of event reports. Lack of knowledge of sensor placement can lead to a false picture of the range and intensity of threats a particular organization faces.
- **The CPP sensor grid does not appear to address encrypted activity.** Many of the DOE's most important services (from interactive login to password-restricted web sites) are often protected by encryption that the CPP sensors cannot analyze very well. The result is that the most important cyber services are the least analyzed.
  - *Recommendation: Create a signature team for the CPP.* The team would provide quality assurance testing and refinement on signatures, identify the most appropriate signatures to deploy in the CPP sensors, and develop signature for newly detected attacks.
  - *Recommendation: Document the DOE site's network architecture and the sensor placements in this architecture.*
  - *Recommendation: Create a taskforce to determine how to analyze attacks that current CPP sensors cannot analyze.* The most obvious cases involve encrypted services, but NAT, internal switched networks, and other trends pose challenges to the current generation of sensors.

## 7 Conclusions

The CPP sensor grid was originally developed to support a counterintelligence mission, and the defensive cyber security mission was added later. Unfortunately, the CPP sensor grid design and current operations poorly support the cyber security mission. In addition to the inadequacy of the CPP sensor grid, the defensive cyber security mission suffers from the lack of a clearly articulated mission that can be accomplished within the expected budget. For example, the budget of a similar program, the Air Force's AFIWC/AFCERT program, is much larger than the DOE's effort, so expecting the DOE to accomplish a similar mission is unreasonable. We recommend the DOE's CPP defensive cyber security effort focus on a strategic intelligence capability.

The report began in Section 2 with a summary of the findings and recommendations found throughout the rest of the report. Next, Section 3 summarized the CPP system as described to us by LLNL/CIAC personnel, supplied unclassified documents, and information available on the web. Section 4 described several actions different people may want to take based on the same information. Our purpose in describing this section was to illustrate that different people have different missions with different goals, so translating information into actionable information will differ depending on who the customer of the information is. Section 5 looked at several different defensive cyber security mission illustrating that different missions place different requirements on their sensors. We spent considerable time in Section 5.4, the strategic intelligence mission that we believe is the best fit for the DOE's CPP defensive cyber security mission. Finally, Section 6 looks specifically at the CPP sensor suite and issues surrounding it with respect to a defensive cyber security mission.

## 8 References

[DOE 00] "Follow-on Review of the Status of the U.S. Department of Energy's Counterintelligence Implementation Plan", DOE/IG-0464, <http://www.ig.doe.gov/pdf/ig0464.pdf>, March 2000.



- [EO 12333] Executive Order 12333, <http://www.cia.gov/cia/information/eo12333.html>, Dec 4, 1981.
- [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 Avoidance", *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.
- [Schr 05] Schroll, Adam. Private Communications, 21 April 2005.
- [Schw 02] Schwartz, John. "Year After 9/11, Cyberspace Door Is Still Ajar," *New York Times*, 9 Sep 2002.
- [Syma 05] Symantec Corp., "Internet Security Threat Report, Volume VII", March 2005.
- [WHO 99] "Principles of Disease Surveillance", World Health Organization (WHO), <http://www.who.int/emc/slideshows/Survintro/sld001.htm>, Oct. 1999.