Global Information Assurance Certification Paper

SANS Security Essentials
GSEC Practical Assignment version 1.2e

Proposed Conceptual Tools for Managing Cost and Complexity When Securing Networks

Kathleen E. Howard

Introduction

Currently a wide selection of network security and intrusion handling methodologies are available and some consistent principles are emerging. Because of the cost and increasing complexity of providing network security, I believe a fruitful and necessary next step is development of a set of conceptual tools to help decide when and how best to invoke these step-by-step measures. This paper will describe the cost and complexity issues facing security professionals, outline the desired outcome in facing these issues, and finally will suggest initial proposals for reaching those goals.

The Challenge

Cost is an issue for everyone, Dorothy E. Denning cites a 1997 survey by *Infosec News* in which 62 % of the information security professionals responding said budget constraints were a significant obstacle to achieving adequate security and another 20% said it was the single greatest obstacle they faced. [1]

As for complexity, Bruce Schneier tells us, “the future of digital systems is complexity, and complexity is the worst enemy of security.” [2] Complexity is also a pretty big enemy of cost containment as the following information from a recent Honeynet Project illustrates.

The Honeynet Project is a volunteer-staffed research group of security professionals using a network of standard production systems to capture black hat activity. The information is then analyzed and shared. The project also sponsors challenges open to all comers to analyze the data the honeynet has captured in the hopes of increasing knowledge in the field and standardizing forensic methods.[3]

One of the realities the challenges underscore is that as providing security has grown in cost and complexity, the virtual opposite holds true for those who would attack systems. Because of the wide and easy availability of hacker scripts and access to computers at home, work and public sites such as libraries, significant damage can be done by someone who has little or no actual computer skill, has invested little or no money, and who has not dedicated much more time than it takes to download and run a script. As Dave Dittrich point out in discussing investigation costs for the Forensic Challenge from earlier this year: “One thing is for certain. It is ***much*** harder and takes more

skill to figure out what was damaged than to do the damage." [4] So much harder that most entrants finished the challenge "when they ran out of time, not when they felt they were done. " [5]

As a by-product of this effort, information about the cost and complexity of such forensic investigations is also being collected. For the challenge mentioned above:

> The average time spent in investigation turned out to be about 34 hours per person. That's a standard week's worth of work to clean up and deal with the mess left by an intruder in about a half an hour. That's about a 60:1 ratio! Using a standard upper-mid range annual salary figure of US$70,000 per investigator, that works out to be a cleanup cost of over US$2000 for a single incident. It is very likely one of dozens, if not hundreds, of intrusions just like it. [6]

Dittrich conclude this "is not to suggest that ***every*** intrusion warrants a complete forensic investigation, but in some circumstances it is entirely appropriate and needs to be done quickly (and correctly)." [7]

Please note, this challenge was purely analytical, it did not include the time it would take to actually fix the affected systems, let alone the time it would take to deal with the evidence collection process if an organization decided to prosecute. It is unlikely there is any individual or organization which could apply the level of effort described in most standard incident handling methodologies to every single piece of unexplained traffic detected. A survey of just three step-by-step approaches available on the internet (the Navy's "Computer Incident Response Guidebook," [8] an "Operations Manual -Information Protection Center" [9], available at SecurityFocus.com and a tutorial on investigating computer crime from the same site [10, 11]) all provide representative (and extensive) detail on what to do when you get an incident, the importance of a standard methodology and examples of reports, checklists, and procedures. What is missing is a context for overburdened and under-funded security staff to use in making informed decisions on when to apply all these steps.

So, how do security professionals even the odds a bit? What rules of thumb can be applied to provide a sensible cost/benefits analysis in deciding whether to invoke the standard forensic methodologies and at what level of detail?

The Goal

With these challenges in mind, the goal is to provide the greatest protection for the least cost in hardware, software, bandwidth and human effort. Human effort includes not only the time needed to investigate and recover from an incident, but also the level of expertise/training and skill required.

One of the keys to managing costs is simply to make it an explicit part of the decision making process. It isn't always easy to quantify the cost of an incident

or the value of a preventive measure, but collecting metrics and comparing results to those of other organizations is the first step in standardizing this aspect of computer security field.

Managing complexity is the second major focus in leveling the playing field between attackers and defenders, and it assists in managing costs in the process. Several facts provide the foundation for this approach:

- Not all incidents are equal (a port scan is not equivalent to a rooted server)
- Therefore, not all incidents require the same level of effort to manage.
- 100% security is probably not feasible in most environments and cost is likely to be the determining factor in how safe you can be.
- Limiting complexity will generally limit cost.
- Limiting data can play a significant role in limiting complexity in the forensics field.

A Proposed Model

These facts translate into several rules of thumb that can be brought into play to limit complexity, maximize resources and limit risk:

- Stop malicious/suspicious traffic as close to the perimeter as possible.
- Match the protection effort/expenditure to the threat level.
- Use the least costly protection feasible to manage each threat type.
- Limit analysis to the minimum data needed to justify/support your actions.

All well and good, but how do you do it? Below are some conceptual tools to help provide a context for making intrusion handling decisions and assist in managing the scope of activity by managing the complexity.

A context for decision-making:

Much of the literature on network and information security uses a warfare analogy. At least at a superficial level a warfare analogy assumes a winner, a loser and an end, not to mention life-and-death stakes:

- A warfare paradigm presumes a zero-sum winner-takes-all outcome. Because intrusions may never be detected and because the motivation of the intruder can range from curiosity, and a desire to alert the owner to vulnerabilities, all the way up to a government-supported attempt to deny service or steal or corrupt information, it is hard to fit all the actors into such a simple win-lose dichotomy.
- A warfare analogy presumes an end. Unless you live in the Middle East, Northern Ireland or certain parts of Eastern Europe, there is an assumption that the “conflict” will be finite not on-going.
- This approach also assumes a life-and-death ultimate cost justification. Lives of those in the country and the existence of the country itself are at stake. While sufficiently malicious hacking could kill a box utterly, most can be recovered and the same is true at the organizational level, the

payoff isn't killing, it is gaining information, bragging rights, commercial advantage or money, resources or data. Granted, data and resource control could be used to aid in killing a country or company's infrastructure and that could lead to human death.

I agree with Schneier that a better analogy in which to frame this situation is one of risk management, security from an insurance perspective. [12] Not as sexy, but it puts cost directly into the formula. You are not eliminating the "enemy" you are managing risk. Another advantage of this perspective is a focus on process rather than end-game. Businesses and individuals accept the need for insurance for just this reason, bad things, from a variety of sources (health challenges, accidents, crime, natural disasters) happen to individuals and organizations on a regular enough basis to warrant a fallback position - an insurance policy to mitigate the damage, providing assets to aid in the recovery. It is not possible to be completely safe, that is not a war that can be won, and therefore, arguably, it is not a war at all.

It also puts a more realistic focus on the challenge. You are not likely to lose your business due to a hack attack, nor are you likely to kill the intruder. You aren't vanquishing an enemy or being vanquished, you are trying to protect your assets and carry out your organization's business with as little interference as you can reasonably afford. Probes and attacks against your networks, hardware and data are simply another risk you insure against along with theft, accidents, natural disasters and health challenges affecting your employees.

For now, securing your networks from intrusion and corruption is pretty much a do-it-yourself proposition. But a perspective that focuses on matching cost to risk provides a very helpful guideline at every step of the process.

Keeping it simple:

Another useful concept applicable at every step is one borrowed from the object-oriented programming world and popularized there by Grady Booch. That is the theory of information hiding or levels of abstraction. [13] Because object-oriented programming is so complex Booch applied the concept of functional decomposition, only focusing on the information that is relevant at a given level of abstraction and hiding the rest. In essence, increased detail equals greater cost in time, bandwidth, hardware storage, CPU capacity, software sophistication and human intervention. So finding ways to limit the data that must be analyzed both simplifies and expedites the decision making process and limits cost.

Finally, a couple simple time-management truisms are extremely helpful. "Begin with the end in mind," [14] and match effort to the priority of the task as Steven Covey exhorts.[15] There are really only about a half dozen realistic outcomes/actions for any investigation:

- Adjust the perimeter protection device configuration (block IPs, ports,

protocols).
- Harden vulnerable devices (patch or upgrade operating systems, tighten passwords, change file protections).
- Recover compromised devices.
- Restore data.
- Document investigation results (including false-positive findings) if warranted.
- Protect and collect chain of evidence if prosecution is warranted/desired.

It is also relatively straightforward to devise a standard set of categories for intrusions based on the different actions required for each:

- The first step in the triage of any event is to determine whether the detection is a false-positive, most of which are likely to be handled in the same manner.
- If not, then identify if it is a probe, a denial of service attempt, an unauthorized access attempt, etc. You may want further categories for successful and unsuccessful attempts, for different types of access attempts etc. The guideline is to limit the categories to groups that require a distinct set of handling instructions.

So, lets put these concepts into play in a couple examples with the SANS defense-in-depth approach.

Our first step would be to look at the highest level of abstraction and resolve as many security issues as we could at this level. At this level we basically have a very simple us vs. them model. Our network is viewed as a single entity and the outside world is divided into legitimate and non-legitimate traffic.

At every level of abstraction we will keep our insurance analogy in mind: Is the cost of this protection reasonable in light of the benefits it provides? Is there a way to limit cost and/or increase benefits?

Outcomes: Adjusting perimeter protection devices to drop as much invalid traffic as possible, as automatically as possible, while maintaining full functionality for users/customers and possibly documenting results (via log collection and analysis) are our only likely outcomes.
Question: The question we are trying to answer is: At this level what can we do to identify and block invalid traffic and permit valid traffic?
Data:

- The traffic data we have available at this level may include source and destination IP, ports, protocols, time and packet size.
- Another source of data is vendor, security community information, and other research about ports and protocols (and even IP ranges) it might be advisable to block.
- Valid internal IP ranges, ports and protocols required for valid operations.

Costs/Efforts vs. Benefits/Priorities trade-offs:

- Sufficient research to determine what ports are likely to be exploited on the one hand and which are required by users or customers to conduct business on the other hand. Similar trade-offs must be examined for IP blocks and protocol blocks. This is an ongoing and non-trivial task as work is being done all the time to find new and different ways to exploit commonly used communications patterns (see Ofir Arkin's paper on using crafted ICMP packets to map not only a network, but many common operating systems [16]).
- Another decision at this level is how much standard logging to do. Again it is a matter of striking a balance between the cost of collecting, storing, and at least occasionally analyzing data vs. the cost of not having data on-hand to assist in an investigation.

For traffic that gets inside the perimeter but triggers an intrusion detection system (IDS) signature, the same analysis applies:
Outcomes: We match our action to the category of event the incident falls into (this could include a block or blocks at the perimeter level, hardening a box or recovering a box as well as documenting the action and/or collecting data for prosecution).
Question: The question we are trying to answer is: What category of event are we dealing with and what actions match that category?
Data:

- The traffic data we have available at this level may include source and destination IP, ports, protocols, time and, depending on the IDS, raw data, packet information.
- Another source of data is destination IP information: is it an IP in use, what is the operating system, what version is it running?

Costs/Efforts vs. Benefits/Priorities trade-offs: The categories of attack would need to take into account cost vs. benefits in the following questions:

- Is the target vulnerable to the attack? We may want to block the source regardless, if we have the resources and feel it was a malicious attempt.
- Was the target compromised?
- What is the priority of the attack? Depending on volume and organizational policy we may not block any probes, or attacks aimed at non-vulnerable machines. Successful and unsuccessful attempts might be handled differently.

At this point a very real danger is getting lost in, or overwhelmed, by the data. A recent GAO report on Critical Infrastructure Protection noted the organization being evaluated "did not have computers capable of rapidly analyzing the large amounts of data associated with some cases." [17] Once the amount, complexity and number of sources of data reach a certain level, it isn't just hardware constraints that become a problem. It is also an issue of the capacity and sophistication of the software and the skills of the individuals doing the analysis.

Given the time, manpower, computing and bandwidth constraints faced by virtually all organizations, a big part of the job at this level is simply deciding where to draw the line on analysis and response. In a world with unlimited resources it would be possible to block every unauthorized probe, no matter how limited, to launch a forensic investigation of every suspicious IDS event, no matter how isolated, or how invulnerable its target.

In the real world this is where cost becomes a major determining factor in deciding what action if any to apply.

Clearly, the most resources must be applied to identify, and counteract attacks which successfully target and compromise vulnerable boxes. This means preventing further access to your network from the offending source (not as simple as it sounds in this world of spoofed IP s and ISPs with DHCP). Evaluating and repairing any damage to the box or boxes compromised and, depending on the policy of your organization and the severity of the compromise, securing the box and collecting data for prosecution of the intruder.

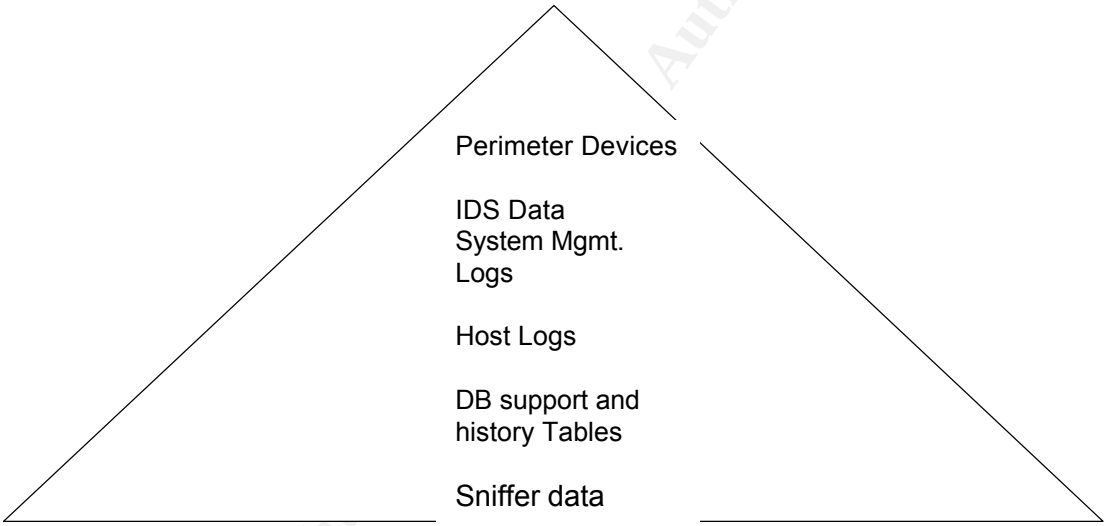

Beyond this obvious approach, any investigation which involves collection and cross-correlation of multiple data sources, significant quantities of data, coordination with multiple individuals, organizations or hardware or software systems, will benefit greatly by staying as close to the top of the data pyramid as possible. For every step down into a more complex level or greater quantity of data, scope the effort by considering what the end result is likely to be, what the priority is, what the minimum level of data is to justify/support the end result, and whether the cost of obtaining and evaluating that data and obtaining that end result is justified in light of the benefits it will provide.

Conclusion

There are several tests of the value of a model. Does it make a complex

conceptual idea easier to discuss and visualize? Does it make it easier to implement the idea in the real world? Does it provide a way to measure success? Does it make it easier to see where the gaps in technology and methodology are?
I am hopeful this discussion at least met the final criteria and we will soon see some comprehensive models to deal with the ever-increasing spiral of cost and complexity in the network security field.

List of References

1. Denning, Dorothy E. Information Warfare and Security. Reading: Addison Wesley Longman Inc., 1999. 396.
2. Schneier, Bruce. Secrets and Lies: Digital Security in a Networked World. New York: John Wiley & Sons, Inc., 2000. 353-354.
3. "About the Project" The Honeynet Project URL: http://project.honeynet.org/project.html (4 July 2001)
4. Dittrich, Dave. "Results of the Forensic Challenge" The Honeynet Project URL: http://project.honeynet.org/challenge/results/index.html (4 July 2001)
5. ibid.
6. ibid.
7. ibid.
8. Department of the Navy, "Computer Incident Response Guidebook Module 19 Information Systems Security (INFOSEC) Program Guidelines," August 1996. URL: http://www.nswc.navy.mil/ISSEC/Docs/P5239-19.html (4 July 2001)
9. Mackie, Andrew. "Operations Manual Information Protection Centre" 28 April. 2000 URL: http://www.securityfocus.com/focus/ih/ipc/OPS-man-admin.htm (4 July 2001)
10. Wright, Timothy E. "An Introduction to the Field Guide for Investigating Computer Crime." 17 April 2000. URL: http://www.securityfocus.com/frames/?focus=ih&content=/focus/ih/articles/crimeguide1.html (4 July 2001)
11. Wright, Timothy E. "The Field Guide for Investigating Computer Crime: Overview of a Methodology for the Application of Computer Forensics Part 2." 26 May 2000 URL: http://www.securityfocus.com/frames/?focus=ih&content=/focus/ih/articles/crimeguide2.html (4 July 2001)
12. Schneier, Bruce. Secrets and Lies: Digital Security in a Networked World. New York: John Wiley & Sons, Inc., 2000. 383-386.
13. "Information Hiding" URL: http://www.webopedia.com/TERM/i/information_hiding.html (4 July 2001)
14. Covey, Steven R. The Seven Habits of Highly Effective People. New York: Simon and Schuster, 1989. 98.
15. ibid. 150-154.
16. Arkin, Ofir, "ICMP Usage in Scanning." V. 2.5 Dec. 2000 URL: http://www.sys-security.com/archive/papers/ICMP_Scanning_v2.5.pdf (4 July 2001)
17. U.S. General Accounting Office. "Critical Infrastructure Protection Significant Challenges in Developing National Capabilities." GAO-01-323 April 2001. 65. (URL: http://www.gao..gov)