

TRAINING THE CYBER WARRIOR

J.D. Fulp

Naval Postgraduate School

Abstract: This paper suggests the major educational components of a curriculum that is designed to educate individuals for job assignments as Information Assurance professionals – also known as: cyber warriors. It suggests a minimum common body of knowledge for all cyber warriors along with two major specialization categories: cyber tacticians and cyber strategists. The paper describes the distinction between tactician and strategist and offers a rough outline of the education each should receive.

Key words: Education, Cyber Defense Exercise, Information Assurance Curriculum, Cyber Tactician, Cyber Strategist, Cyber Warrior

1. INTRODUCTION

Though the wide scale interconnection of automated information systems (e.g., the Internet) has been a boon to U.S. military and economic power, it also presents a soft underbelly to present and future adversaries. U.S. reliance upon the ever-expanding National Information Infrastructure, in conjunction with an increasingly wired world, exacerbates the U.S.'s vulnerability to cyber threats. Sensitive information could once be protected using relatively easily understood physical, personnel, and communications security mechanisms. The advent of interconnected automated systems now requires that such information receive the additional protections afforded by computer and network security mechanisms. The education of individuals to understand the complexities inherent in such mechanisms, so that they can effectively implement them is the central theme of this paper.

It is natural to invoke military principles and terminology to discuss elements of this new era of increased cyber vulnerability. Though those involved may not wear uniforms, or fight along linear geographic

boundaries; there is clearly a high stakes adversarial environment that is conducive to well understood military conceptualization. Therefore we have such analogies as: de-militarized zones (moderately protected public service networks), security perimeters (boundaries between different data risk levels), and cyber warriors (IA professionals). This paper suggests an education regimen for cyber warriors, and suggests that these warriors be divided into two categories: cyber tacticians and cyber strategists. Cyber tacticians would focus on reducing the risk of existing fielded systems primarily through the application of appropriate safeguards (e.g., firewalls, intrusion detection, redundant configurations, data backups, etc.). Cyber strategists would focus on reducing the risk of future systems primarily through the application of structured and formal system design techniques that reduce system vulnerabilities.

2. BOOT CAMP: TEACHING THE FUNDAMENTALS

Cyber tacticians and cyber strategists should both receive the same basic core education. We can refer to this as cyber boot camp in keeping with the military analogy. Cyber boot camp should address all of the core subject matter encountered in modern information systems, and do so in a bottom-up order. Cyber boot camp should also introduce the core principles of information assurance.

2.1 Subject Matter Ordering

The minimum set of core subject matter courses I suggest are: 1) Discrete Mathematics, 2) Computer Hardware/Architecture, 3) Programming, 4) Operating Systems, and 5) Algorithms. These choices will be elaborated upon below. By bottom-up order, I suggest introducing the courses in the order presented above while allowing that Programming and Operating Systems may be presented in any order due to their logical interdependency. Teaching these courses in the order suggested should reduce much of the confusion often experienced by novice students who find themselves working with an abstract logical concept (e.g., pointers) before they have seen the underlying physical level implementation (e.g., a 32 bit address indicating a physical memory location). Upon completion of this course of study, students should have a clear understanding of the problem-to-solution process in its entirety. That is; real-world problem statement → algorithm to solve it → program to implement the algorithm → operating

system that will load, schedule, and allocate resources for the program → hardware that electrically executes the instruction-cycle and runs the loaded program that solves the problem. And this entire process can be conceptually described or modeled with the most basic layer, i.e., Discrete Mathematics. Throughout the presentation of these five core courses, we also introduce the students to the core principles of information assurance that are described in section 2.3.

2.2 Subject Matter: Courses

Cyber warriors should begin with the solid conceptual understanding that computers are ultimately nothing but physical structures that provide a means for mathematics to be brought to corporeal life. For example, numbers translate into; pixilated screen images, hard drive armature displacements, pointer offsets, IP address masks, etc. Discrete Mathematics provides the descriptive tools necessary to discuss, define, design, and analyze the behavior of computer hardware and software. It is the logical starting place for the study of information processing systems, and provides the necessary tools for describing system design and functionality.

Next in the sequence is computer hardware and architectural design. At this layer students learn about the binary switch (transistor) that is at the “atomic level” of computer processors. They learn about logic gates that realize Boolean relationships that make the controlled movement and manipulation of digital data possible. They are introduced to combinational circuits, storage devices, encoders, decoders, multiplexers and the other basic digital building block components. The capstone instruction in this course should consist of a demonstration of how a high level language code fragment must be converted to machine code that is supported by the underlying target hardware’s instruction set; followed by a clock-cycle-by-clock-cycle analysis of what happens as the hardware processes each fetched instruction from memory.

We follow the hardware layer with a course in either Operating Systems (OS) or a contemporary programming language. On the one hand, OSs are programs, thus we would expect programming to precede the OS course. On the other hand, application programs rely upon an appropriate OS environment upon which to run. This classic “chicken or egg” relationship should not be a major point of contention as it pertains to the quality of our cyber warrior curriculum. As indicated above, I suggest that instruction in programming follow directly behind the hardware course. This allows students to immediately “use” their newly understood knowledge of hardware by writing instructions that will ultimately run on it.

With the rudiments of hardware, and software design understood, we should next instruct our students in the features and functionalities of operating systems (OS). Students come to understand the central role that the OS plays in choreographing the interaction between special purpose application programs and the host platform (hardware) it is being run on. It is also at this point that our students should begin to see how a well designed OS can play a crucial role in a cyber defense strategy via such mechanisms as: file system support for access control, subject/object labeling, locking mechanisms, security domains, and segmentation.

We finish the coursework with our students being indoctrinated into the world of complex problem solving with a course in advanced algorithms. Students learn that size and speed matter in computer systems just as they do on the battlefield. They also obtain enhanced understanding of the complexity of the myriad protocols employed to bind systems in an inter-operative networked environment. Our cyber warriors are now mentally armed to scrutinize the complexities of such topics as: key space search efficiency, path finding, tree pruning, shortest path determination, signature matching, etc.

2.3 Core Principles of Information Assurance

Throughout boot camp we inculcate the students with the seminal concepts in information assurance methodology. Though many concepts, principles, models or theoretical postulations may legitimately vie for inclusion in this category, I suggest the following four as an absolute must: the Reference Monitor Concept, the Risk Management Equation, the Defense-in-Depth paradigm, and the Principle of Least Privilege. The Reference Monitor Concept is at the heart of virtually every technical mechanism (hardware or software) that has ever been devised for the purpose of enhancing the security of information. The Risk Management Equation provides a high level management framework by which cyber warriors can organize and allocate their defensive efforts. Defense-in-Depth dictates not putting all of one's security eggs in one basket, but instead employing multiple, sometimes overlapping, layers of complementary security solutions. The Principle of Least Privilege enjoins all who develop or configure security-relevant attributes of systems to allow no more access to information or computing resources than is absolutely necessary to accomplish each legitimate (i.e., non policy violating) task.

2.3.1 The Reference Monitor Concept

The Reference Monitor (RM) Concept, first introduced in the “Anderson Report” [1], provides the most basic and essential technical framework for any information assurance solution. I will make no attempt at a complete description here, but I will offer a synopsis that highlights the importance of this concept to the proper education of the cyber warrior.

The concept maintains that access control is at the heart of data protection. An access request is defined as a subject (person or process) attempting to read or modify an object (logical unit of data). The RM is the mechanism that arbitrates such requests, and does so based upon one or more identifiable or otherwise measurable attributes associated with each subject and object. The actual access control policy that a given RM implements is determined by the relationship of the subject and object attributes and the rules that the RM enforces over these relationships. A more thorough examination of this concept can be found in “The Reference Monitor Concept as a Unifying Principle in Computer Security Education” [2].

2.3.2 The Risk Management Equation

The Risk Management Equation gives the cyber warrior a big picture management perspective over the extensive problem domain of IA. The equation is derived from the generally accepted notion that safeguards applied to mitigate initial risk will reduce that risk to some degree, resulting in residual risk. This can be expressed relationally as: zero risk \leq residual risk $<$ initial risk, and from that the more general relationship is shown:

Residual Risk = Risk - Safeguards

Then, applying the notion put forth by Brinkley and Schell [3], we can substitute the product of threats and vulnerabilities (abbreviated “Vulns” below) for risk, to achieve the final risk management equation.

Residual Risk = (Threats x Vulns) – Safeguards

When explaining this equation, students must be informed that merely defining information as the subject of this equation yields insufficient granularity. Instead, students learn that there are ultimately four attributes of information that are potentially of interest to protect: confidentiality, integrity, authenticity, and availability. These attributes are so central to IA that we exhort our warriors to always be mindful of these four attributes when investigating any given IA question. We should refer to these attributes so often that the acronym CIAA becomes part of the cyber warrior’s lexicon. Note that while some IA practitioners recommend adding non-repudiation to the list of protected information attributes, I recommend

omitting it as it is essentially a byproduct of sufficiently implemented authenticity along with integrity of a one or more attendant timestamps.

Now that our students understand that the Risk Management Equation can be defined collectively over all information attributes, or more precisely, over any of the four specific information attributes, we can proceed to discussion of the equation’s individual terms and their relationship.

The threat vulnerability product is somewhat intuitive, but deserves a brief explanation for our novice students. Threats indicate malicious intent to attack one or more of the four information attributes. Brinkley and Schell [3] describe six such threats: human error, abuse of privilege, direct probing, probing with malicious software, direct penetration, and subversion of security mechanism. Vulnerabilities indicate design flaws in the security mechanisms of a system. The product of threats and vulnerabilities is equivalent to risk. Expressing risk as the product of threats and vulnerabilities captures the logical conclusion that risk does not exist for systems that have no vulnerabilities, and conversely, that the lack of any threats poses no risk no matter how many vulnerabilities a system may have. Increasing or decreasing either of the two product terms yields a corresponding increase or decrease in risk. To mitigate risk we apply safeguards, which if effective, should reduce the risk by some amount leaving us with residual risk. Since it is generally considered infeasible to achieve a zero residual risk environment, our cyber warriors are taught that their broad mission is to manage the equations’ three dependent variables (threats, vulnerabilities, and safeguards) in such a way as to reduce residual risk to an economically (or militarily) acceptable level. This understanding yields a simple big picture IA management matrix.

	Threats	Vulnerabilities	Safeguards
Confidentiality			
Integrity			
Authenticity			
Availability			

This matrix identifies twelve areas of concern to the cyber warrior. We can teach this as a mental model that the cyber warriors can use in their daily routine. For example, cyber warriors make the checking of new vulnerability alerts (e.g., CERT advisories) a part of their daily routine. As new

vulnerabilities are discovered and announced, they are quick to assess which attribute(s) of information the vulnerability applies to, and what the resulting impact will be on the residual risk of information under their protection. Likewise, these cyber warriors will monitor developments among IA vendors for improved safeguards, ready to investigate and perhaps recommend for purchase any products that promise a reduced residual risk return on investment.

2.3.3 Defense-in-depth

This core principle dictates that practitioners of IA should not rely on any single device, technology, or security area (e.g., personnel security, physical security, etc.) when working to minimize system risks. Practitioners should instead seek to bolster system defenses by incorporating multiple devices, technologies, and security areas in a synergistic and mutually supportive manner. The new Department of Defense Directive on Information Assurance [4] emphasizes the importance of this core principle by addressing it in its very first paragraph:

1.1. Establishes policy and assigns responsibilities under reference (a) to achieve Department of Defense (DoD) information assurance (IA) through a defense-in-depth approach that integrates the capabilities of personnel, operations, and technology, and supports the evolution of network centric warfare.

The DoD IA Directive also provides a definition for this core principle:

E2.1.11. Defense-in-Depth. The DoD approach for establishing an adequate IA posture in a shared-risk environment that allows for shared mitigation through: the integration of people, technology, and operations; the layering of IA solutions within and among IT assets; and the selection of IA solutions based on their relative level of robustness.

We should emphasize the importance of defense-in-depth with case studies where a seemingly sufficient single layer defense proved insufficient. We should stress to the fledgling cyber warriors the extreme skill and resolve that some attackers will bring to bear in a concerted assault, and the value that a layered defense-in-depth approach provides in countering such attackers.

2.3.4 Principle of Least Privilege

The Principle of Least Privilege is the more all encompassing principle that borrows directly from the intelligence community's institutionalized "need to know" personnel security principle. The idea is that sensitive information should receive no more exposure to potential disclosure or modification risk than which is absolutely necessary for mission accomplishment. Cyber warriors should be taught to employ the least privilege principle to the maximum extent practical; including: user account privileges, listening ports on servers, ICMP response messages (e.g., no response to echo request or requests for subnet mask information), and firewall permit rules, to name a few instances. Least privilege is equally applicable in software design where, for example, we would expect the operating system to restrict a given module's instruction space (i.e., branching) to that module's assigned/allocated memory segment; or via the use of "friend" class relationships in object-oriented programming to restrict illicit or otherwise erroneous inter-object message passing (member function calls).

3. TWO CATEGORIES OF CYBER WARRIORS

Though it is possible to educate all cyber warriors the same, the breadth of the IA problem domain coupled with economic realities suggest specialization should be more granular. I suggest that the two top-level categories be: cyber tactician, and cyber strategist. The military analogy is strong here but not exact. Cyber tacticians focus on reducing residual risk predominantly with the application of safeguards, while cyber strategists focus on reducing residual risk by reducing system vulnerabilities. The skill set for each of these risk mitigation solutions is sufficiently different and complex to warrant specialization.

3.1 Cyber Tacticians

Basically, cyber tacticians should be educated to protect the systems that are fielded now. Due to economic forces and improperly educated or motivated programmers, computer systems will be fielded with vulnerabilities that run the gamut of type and severity. Since history gives us no hint that threats are subsiding, the risk management equation tells us that these systems are at risk and that the application of safeguards is the only in-field means of risk mitigation.

We should educate the cyber tacticians to become experts in the utility, application and effectiveness of safeguards. An exhaustive list of safeguards is not attempted here, but the broad categories of safeguard tools and technology are.

First on this list is the important though mundane category of secure standard operating procedures and user training. This category covers such items as: password selection and usage, un-attended log-ons, potentially malicious e-mail attachments, the importance of anti-virus signature updates, social engineering attacks, portable PC security, etc.

Second is data backup technology and policy. The efficacy of this safeguard category is widely known, but the confusion encountered by the multitude of media (e.g., tape, disk, CD-R/RW, DVD+R/RW/RAM) and backup techniques (e.g., full, incremental, differential, RAID levels 1-5, compression, etc.) dissuade many well intentioned users from making it a part of their routine data security habits. Cyber tacticians should maintain mastery of this extremely important recovery safeguard technology, and ensure that a backup policy is created and implemented for all valuable data under their purview.

Cyber tacticians should be well versed in the “principle of least privilege” as it pertains to all aspects of information security. This principle should permeate every configurable software setting and every access control decision. Cyber tacticians should know that vendors often practice the “principle of most privilege” as their out-of-the-box default configurations, including generic root logins and passwords. A regular and concerted effort to ensure that a least privilege policy is enforced system-wide should be heavily stressed.

We should educate the cyber tactician to make regular checks for newly announced vulnerabilities, and be proactive in seeking and installing vendor patches as soon as they become available. The tactician should also be able to assess the added risk that any announced vulnerability presents, and be prepared to take other defensive measures until a patch is available. The measures might include a modified firewall rule-set, proxy isolation of the vulnerable service, or even removing the service from the network in extreme cases.

Cyber tacticians should learn the value of redundancy for systems, services and power. They should be taught to assess an agency’s high value data or service assets and be able to propose, design, and implement a redundant/failover configuration that enhances data and service survivability.

We should teach our tacticians to do regular vulnerability assessments of their own systems, thereby taking a proactive role in identifying defensive weaknesses before an attacker does.

Cyber tacticians should be educated as experts in choosing and configuring firewalls and intrusion detection devices and software. They should learn the types of filtering (e.g., stateless, stateful, reflexive, proxy-level, etc.), how to understand and build filter rule-sets, and how to interpret packet level information (e.g., TCP flags, TTL values, sequence numbers, etc.). These skills enable the tactician to read and understand network traffic logs, identify anomalies, and react to such anomalies with updated filter rules.

Encryption technology is next on the agenda. Cyber tacticians should know the fundamentals and ramifications of such cryptological concepts as: block versus stream ciphers, chaining, key symmetry, key space, key management, hashing, and common protocols used to implement secure network transactions (e.g., ISAKMP, IKE, SSL, SSH, IPsec, PPTP, etc.). Cyber tacticians should be capable of configuring appropriately secure communication tunnels between any two protected systems.

We should teach tacticians the art and science of post-incident computer forensics so that they can sift through the digital residue left in the wake of an attack. They should learn how information is stored and how it may be deliberately hidden or subverted. They should learn the tools and techniques of logging, disk examination, evidence recovery, and legal preparation.

Finally, we should complete the cyber tactician's education with several practical cyber defense exercises. These exercises would entail the design, installation and configuration of a highly secured service network. This network would then be the target of attack by a team of cyber warrior graduates who would employ their knowledge and all available exploit tools to try and compromise the protected network. The earlier attacks can be escalatory in nature so that the defending students can more easily observe and learn. For example; the attackers would first engage exclusively in reconnaissance or discovery type activity (e.g., foot printing, port scanning, etc.), followed by surreptitious attacks intended to achieve unnoticed account access or observation of data, then attacks that modify data, and finally the more brutish denial of service category of attacks. Later exercises should be "free play" for the attackers while the cyber defenders must be on guard for anything. Attack/defend exercises such as this provide realistic scenarios that puts to practice the previously mentioned areas of cyber tactician education. The cyber tactician that has her network: 1) patched, 2) configured for least privilege, 3) scanned for vulnerabilities, 4) monitored by network and host-based intrusion detection systems, 5) properly isolated with proxies and/or firewalls, 6) backed up, 7) redundantly configured, and is herself capable of 8) forensic analysis; has vastly minimized her network's residual risk with sound defense-in-depth IA safeguards.

Students of IA at several of the Service Academies and the Naval Postgraduate School (NPS) have participated in two such large scale exercises since 2001 [5]. In these two exercises, the IA students at each school (Blue Teams) configured nearly identical service networks, and applied to these networks the security principles learned in their IA courses. A Red Team comprised of information warfare professionals from the NSA, Air Force, and Army, then attacked each network through a VPN tunnel for four consecutive days. Each Blue Team was graded based upon the resilience of its network to attack, and the accuracy of its daily situation reports which identified each day's attack activities, and the success or failure thereof. NPS was the high scoring Blue Team in both of these exercises. A third exercise is scheduled for April of 2003.

3.2 Cyber Strategists

As mentioned above, the focus of the cyber strategist is to reduce risk by reducing system vulnerabilities. By referring back to the risk management equation, we can see that a system with no vulnerabilities results in no risk, thereby negating the need for "after-the-fact" safeguards. The zero vulnerability system is the ideal pursued by cyber strategists, and achievement of this requires a much more theoretical skill set than that of the cyber tactician. So unlike the cyber tactician who builds a virtual protective wall around soft systems, the cyber strategist builds hard systems that need no wall.

Cyber strategists must receive intense education in programming, programming languages, processor functionality, technical policy, and the mathematical skills necessary to understand, code, and formally model the behavior of computer code. This is because the cyber strategist's primary function is to oversee system and network design, development, integration and processor (hardware) functionality to ensure that they correctly implement a given security policy.

Cyber strategists are taught that it is infeasible to attempt to design large general-purpose operating systems to be provably devoid of vulnerabilities due to the arduous and exacting nature of the formal methods methodology required. Instead, they are taught to consolidate all Reference Monitor implementing code into a relatively small software module that is referred to as the security kernel. Strategists must learn the tools and methods by which to ensure the kernel code adheres to three necessary attributes: complete, isolated, and verifiable. A complete kernel is one that is always invoked when any security sensitive access control decision is made. That is, it is proven that no artifice exists that might cause the kernel to miss or otherwise not arbitrate a subject to object access attempt. An isolated

kernel is one that cannot be subverted by any means. For example; booting off of a virus-infected disk, downloading a Trojan horse, or even a human attacker with user level system privilege should not be able to modify the operation of the kernel. A verifiable kernel is one that is small enough to have had every line of code formally proven to be correct. Since the security kernel is essentially considered the first and last line of defense, no chances are taken with its design.

Cyber strategists should study existing systems that satisfy the requirements outlined above, and be presented with instructional security kernel fragments that test their ability to find flaws or prove correctness.

Finally, we should teach the cyber strategists the process and methodology of performing Certification and Accreditation (C&A) so that they may utilize their analysis skills to not only oversee the design and development of new systems, but be able to assess the threats and residual risks associated with existing information processing sites, and be able to make an informed yes/no accreditation decision.

4. CONCLUSIONS

The need for purposefully educated IA professionals is real, urgent, and not expected to abate in any foreseeable technological future. All enterprises with a stake in the protection of information and information processing resources require a knowledgeable staff of cyber warriors to provide it. For maximum return on education investment, cyber warriors should receive extensive education in the following five courses: Discrete Mathematics, Computer Hardware/Architecture, Programming, Operating Systems, and Algorithms. They should receive inculcation in the core IA security principles; specifically, the Reference Monitor Concept, the Risk Management Equation, the Defense-in-depth concept, and the Principle of Least Privilege. Cyber warriors should then select to specialize as either cyber tacticians who focus on the application of safeguards to vulnerable systems, or cyber strategists who focus on the reduction of system vulnerabilities. Cyber tactician education will be steeped in: scanning, patching, least-privilege configuration, perimeter security with filtering, intrusion detection, backup/recovery technology, system/service redundancies, and forensics. Cyber tacticians should participate in several cyber defense exercises to put all of their skills to practical test. Cyber strategist education should be steeped in: formal methods analysis, programming, programming languages, and the tools of mathematical proofing. Cyber strategists should study the design and integration of secure systems and use this knowledge to design future secure systems. Cyber

strategists should also be educated as Accreditors, with a thorough understanding of the complete Certification and Accreditation process and methodology.

REFERENCES

1. Anderson, J. P., *Computer Security Technology Planning Study*. Technical Report ESD-TR-73-51, Air Force Electronic Systems Division, Hanscom AFB, Bedford, MA, 1972.
2. Irvine, C. E., The Reference Monitor Concept as a Unifying Principle in Computer Security Education. *In Proceedings of the IFIP TC11 WG 11.8 First World Conference on Information Security Education*, pp. 27-37, Kista, Sweden, June 1999.
3. Donald L. Brinkley and Roger R. Schell, What is There to Worry About? *An Introduction to the Computer Security Problem*, *In Information Security: An Integrated Collection of Essays*, IEEE Computer Society Press, Los Alamitos, CA USA, 1995
4. Assistant Secretary of Defense for Command, Control, Communications and Intelligence (ASD C3I), *Department of Defense Directive 8500.1, Information Assurance (IA)*, October 2002.
5. Donald Welch, Daniel Ragsdale, Wayne Schepens, *Training for Information Assurance*. IEEE Computer, Volume 35, Number 4, pp. 30-37, April 2002