Proceedings of the 12th Colloquium for Information Systems Security Education
University of Texas Dallas
Dallas, Texas June 2-4 2008

On Developing Tomorrow's "Cyber Warriors"

David Dittrich, University of Washington

Abstract *– Threats of cyber-warfare attacks (and counter attacks) by countries with the largest economies in the world, massive losses of financial and personal data on millions of Americans to cybercrime, and the potential to disrupt Americas critical infrastructures, should be on the minds of all Americans. Why? Because those who design, build, operate and defend the computer systems and networks that our economy relies upon are our fellow citizens. But where will these professionals acquire the skills in Computer Network Operations necessary to secure our future?*

Could a creative, modular design of IA-specific topics allow an educational institution to increase the number of CNSS elements mappable to an undergraduate program, and simultaneously adding "hands-on" learning opportunities to students? Could this model set the stage for expanding CNO education to other programs within the instution, as well as extending partnerships into the broader community? We consider a model that could support delivery of up-to-date demonstrations of current threats found on the Internet, showing students how to protect against, detect, and react to these active threats. In turn, this sets a foundation for establishing a long-term educational path for students that will strengthen the cyber-defenses of our nation in years to come.

Keywords: Cyberconflict, response, expertise.

I. Introduction

In 2003, the Public Broadcast System aired a documentary produced by *Frontline* entitled *Cyberwar* [1] that described attacks on U.S. military and civilian systems going back many years. In 2006, the Chinese news agency *Xinhua* announced a cyberwarfare exercise code named *Vanguard 206B*, just days after the U.S. Air Force announced a summit [2] to bring military operations capacity to cyberspace. Such moves come on top of: the U.S. House of Representatives Committee on Government Reform (which created the Federal Information Security Management Act (FISMA) in 2002) grading nearly one out of three federal agencies it examined as receiving an F in FISMA compliance; The Office of Management and Budget's directive to the heads of all Exective Branch departments and agencies to implement NIST and DoD standard security configurations for all computers running Microsoft Windows operating systems [3]; and the admission from one of the largest chain stores in the United States that a 2+ year long intrusion netted the attackers over 100 million credit card numbers and has cost the company $256 million dollars. [4]

The implication is clear that the U.S. government, civilian critical infrastructure owners, and America's leading businesses need a higher percentage of professionals at all organizational levels who design, implement, manage, and defend Information and Information Systems with Information Assurance (IA) concepts in mind from the start. In addition, they all need to employ (or have access to) a number of highly-skilled individuals that provide the capacity to respond to the largest, most complex, and perhaps most subtle attacks that can be waged against them.

There are perhaps just a few hundred experts in computer security at the highest levels of defensive (let alone offensive) capability in academia and corporations in the U.S. [5] This begs the question, where will a sufficient quantity of highly-skilled CNO professionals come from, and how will they get trained?

Experts in cognitive science quoted in a 2006 Scientific American article entitled, *The Expert Mind* [6] claim that world-class expertise in any given field can be achieved through learning, but becoming expert can require as many as ten years of dedicated practice and constantly challenging one's self against existing experts in the field. In terms of attaining war fighting skills, there are similar time requirements. In an interview on *To the Point*, [7] the former Chief of Plans for American Forces in Bagdad, Lt. Col. Douglas Oliphant, said "We can create a Private in 14 weeks. It takes us 10 years to produce a Major, 20 years to produce a Colonel, and 25 years to produce a General."

It would be laughable to claim that a war could

ISBN 1-933510-96-7 / $15.00 ♡ 2008 CISSE

successfully be fought by Staff Sergeants acting independently, or that it would be possible to take graduates fresh out of flight schools and expect them to fight at "Top Gun" levels of expertise. One can reasonably conclude that it would take more than 10 years of experience at the highest levels of CNO in order to have a cyber defense and offense capacity in line with existing kinetic warfare capacity. So the ideal goal is to have a 10 year path to developing a large base of experts, and simultaneously continuing to retain them and advance them in rank long enough to produce an adequate number of leaders down the ranks to the "battlefield" troops.

The number of major computer security breaches today may indicate we are already a decade behind, even with the U.S. government's efforts to increase the number of trained security experts. While it is a hopeful claim that a CNO expert can be produced in 10 years, there remain many questions to be answered. Is there really a 10-year pathway in today's educational system in which to acquire the skills necessary to counter complex and sophisticated attacks on our critical infrastructure systems by adversarial cyber-warriors? Are today's universities even capable, let alone the ideal place, to produce the large number and high-skill level of information security professionals needed by industry and government to deal with large-scale and concerted cyber attacks being faced *today*? How will we motivate students to stay on a decade-long path to become experts, or how can we accelerate the rate of producing highly skilled CNO professionals within the historically slow-to-change university environment? What other factors will contribute to (or hinder) creating such a sustained 10 year pathway to acquiring expert-level skills?

The National Security Agency's (NSA) Committee on National Security Systems (CNSS) has the responsibility for defining the standard elements of Information Assurance practice for Information Security Professionals, Managers, System Administrators, Auditors, etc.[1] They also have responsibility for certifying that curricula at CAEs map to these standards, and managing the IA Scholarship Program (IASP), which includes scholarship funding for students seeking degrees from CAEs. Part of the intent for the IASP program, as defined in 1998 by Presidential Decision Directive (PDD) 63, [8] was to promote long-term educational opportunities that will produce the professionals listed above to meet the objectives set forth by the President of the United States for security of systems involved in national security, as well as those supporting our nation's critical infrastructure sectors. Since computing is pervasive across all businesses and public sector organizations, so too must IA education extend as widely as possible, embracing interdisciplinary thought and cross-institutional partnerships.

While the CNSS standards are dismissed by some as being "government-specific," this criticism follows from a mis-understanding of the purpose for the standards, or from over-generalization about the applicability of the standards because they include elements related to federal regulations that only apply to the military or federal government (e.g., national policies regarding control of national security information and protection against threats posed by hostile foreign intelligence services.) These criticisms miss one of the most important aspects of IA, which is the use of a rich and broad description of attributes (confidentiality, integrity, availability, encryption, and non-repudiation) and capabilities (protection, detection, and reaction) that is far more useful than the more general term *information security*. While the term *defense in depth* is also used frequently, IA encompasses the concepts of both defense in depth, as well as defense in breadth.

IA, however, is a subset of the skills necessary for the complete spectrum of Computer Network Operations (CNO), which includes Computer Network Defense (CND), Computer Network Attack (CNA), and Computer Network Exploitation (CNE). In this paper, we will use the term CNO to include both attack and defense skills necessary at the expert level, and IA for the foundation skills that are primarily defensive in nature.

II. A potential model for curriculum enhancement

Leading academics have been proposing the integration of security concepts into IT educational curricula, and building a common body of knowledge, for over 10 years. [9, 10] They initiated the long process of changing academic institutions that are notoriously slow to change, and steeped in their own traditions and processes that have no direct relationship to the rapid changes in information technology that businesses and government must adapt to constantly. [11] An even greater challenge is trying to keep up with the rapid advances in malicious soft-

[1] See: `http://www.cnss.gov/instructions.html`

ISBN 1-933510-96-7 / $15.00 ♡ 2008 CISSE

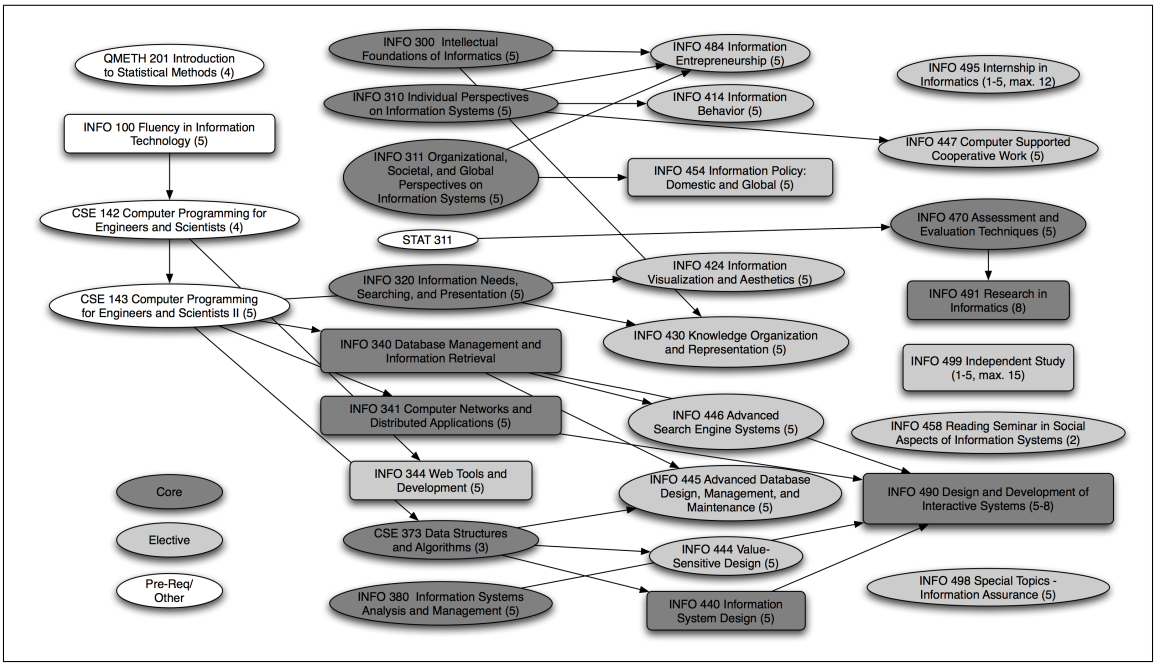

Figure 1: Hypothetical courses w/prerequisite sequencing

ware, created by criminals with financial motives and backing to accelerate their own activities. Effective CNO against today's elite attackers depends on speed and flexibility, which are not fundamental traits of academe. For advanced attack and defense skills, one is hard pressed to find a common body of knowledge that can readily be cited.

This paper proposes a possible way to produce an IA-specific concentration within an undergraduate curriculum, using an existing hypothetical undergraduate degree program as the baseline. We then go further and propose how to integrate these courses with hands-on training in private industry and government, dealing with the actual threats faced by security operations staff in direct response to attacks, as a means of attaining the high level of challenge necessary to create world-class experts in CNO as [6] suggests is possible.

A. A hypothetical undergraduate degree program

Figure 1 depicts the set of courses that make up a possible Bachelor of Science degree program. Courses in dark gray are required courses that make up the "Core" of the program. Courses in light gray are the major electives that the student may chose from that are specific to an undergraduate degree. As shown in Table 1, between 68-73 credits of Core courses, and 12-17 credits of Major electives, are required, with a minimum of 92 credits in all to meet requirements in the undergraduate program.[2]

Pre-requisites	24 credits
Core Courses	68-73 credits
Major Electives	12-17 credits
Undergraduate Program Requirements	92 credits

Table 1: Credit Summary

B. How an IA sub-discipline might be created

A possible path to integrating IA into an existing undergraduate program might be to augment a limited subset of existing courses to include IA modules (shown in Figure 1 as rectangles) and create a very small set of new courses to round out the IA elements in one of the CNSS standards, most likely 4013 (targeted at System Administrators) and/or 4011 (targeted at the *awareness* level for general INFOSEC professional, which includes system developers.) Ex-

[2]The general requirement for a Bachelor's degree might be 180 or more credits, not just the 92 shown in this example. The remaining credits come from liberal studies or other elective courses not directly related to IA topics.

ISBN 1-933510-96-7 / $15.00 ♡ 2008 CISSE

isting courses could be augmented through the use of hands-on modules taught using a Portable Education Network (PEN). [12]

For example, PEN systems support demonstrating network traffic analysis tools in a simulation of what an incident response would be like, using live malware attack tools. The author has prototyped such an lab exercise using a sample of the SSH CRC32 compensation attack detector exploit that was discovered being used *in the wild* in 2001. [13] Feedback from students has been tremendous, as they were able to see what real attack tools look like from the attacker's perspective, from the target system administrator's perspective, and how they can be seen on the network. These kinds of "hands-on" demonstrations engage the students in a way that simple lectures or reading assignments cannot match.

C. *Modular course design and re-use*

Modularity of curriculum below the level of individual courses in an undergraduate degree program is not a common practice. Faculty usually develop all of their own course lectures and labs individually, or in small teams of similarly interested faculty. This results in courses that are cutting edge, and have great depth and relevance to particular sub-disciplines, but at the same time this narrowness of focus can also inhibit the inclusion of concepts – such as IA – that appear to belong to another discipline. At the same time, courses that make up degree programs must be approved by curriculum review committees, and fit within standards set by accreditation bodies.

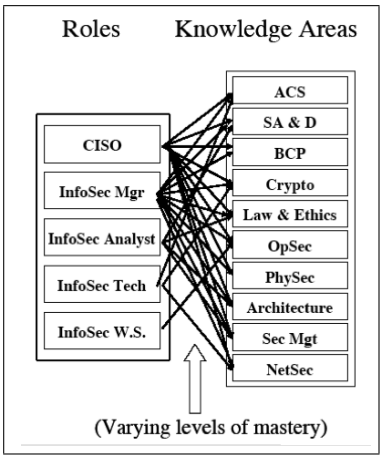

Figure 2: Knowledge area map (Source: [14])

While some argue that these review and accreditation processes, along with other aspects of the higher education system hinder the progress of integration of new material into existing programs, [11] it could also be argued that IA concepts are so integral to topics already being taught that they actually do not constitute "new" material at all, just a different way of teaching information concepts by explicitly addressing *integrity, confidentiality, availability, encryption, non-repudiation*, and the acts of *protecting, detecting, and reacting* to breaches. [10]

Armstrong and Jayaranta [15] show that there is an overlap of skills and tasks that must be performed to bring, for example, evidence of computer crime from the computers and network devices where the evidence exists, to investigators, and eventually to judges and juries. This path involves building a case through the effort of multiple parties, with proper collection, handling and presentation of the evidence in order for it to be useful when the case finally gets to court. See Figure 3.

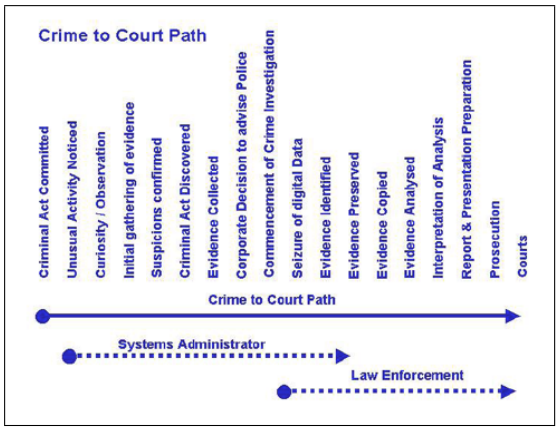

Figure 3: Overlap of Tasks in Crime-to-Court Path of Events (Source: [15])

One implication of this overlap in tasks and skill requirements across different professional disciplines is the need to provide coverage of the required tasks and skills across multiple educational disciplines (in this example, between system administrators and security incident handlers on the technical side, and lawyers, law enforcement agents, prosecutors, and judges on the legal side). The way that modular design and sharing of curriculum can address the overlapping skills requirement is depicted in Figures 4(a) and 4(b).

Figure 4(a) shows how two related topic areas – the key legal issues involved in digital evidence collec-

tion and handling, and concepts of hard disk geometry and file system structures – are almost exactly the same between hypothetical undergraduate degree courses and certificate courses. Producing two separate sets of exactly the same material doubles the amount of work and does not result in an efficiency gain to the institution.

Figure 4(b) shows how the same curriculum materials, if shared with other programs at an institution, could result in more quickly integrating IA modules into existing programs, as well as focusing more energy on higher-level degree programs and building pathways through programs to attract more students and engage them for more lengthy studies.

In addition to commonality of modules horizontally across multiple programs, there is also commonality vertically across differing functional or organizational levels within a given enterprise. As an example, two lectures on rootkits and post-intrusion concealment and log alteration were originally produced at the University of Washington under a National Science Foundation grant supporting a computer forensic program delivered at Highline Community College. These same lectures were later delivered as part of a Special Topics course within the University of Washington's Information School's undergraduate program, and at the University of Idaho in a course on Forensics. This confirms the findings of [14] and [10] that *slip-streaming* through *modules of instruction*, covering issues faced by system administrators, incident responders, and forensic analysts, are for all practical purposes *identical across programs and disciplines* and creation of new modules can benefit multiple courses and programs.

III. Demonstration Framework

The PEN is an isolated network in a box on wheels. Standard builds provide sufficient hardware to simulate a small corporate network. [12] One of the largest time investments in a PEN is in configuration and management of the servers and network devices. Specific skills in administering the network devices, such as Cisco PIX and IOS command line interfaces, are necessary to configure and debug the networking devices, and system administration experience with Windows and Linux are necessary for setting up target servers and attack systems with tools for exploitation. Also, experience in TCP/IP networking and using network monitoring tools are necessary to both use and demonstrate exploits and other security related tools. This goes well beyond the expertise of most faculty who are not spending the bulk of their time teaching network security courses. Understanding the details of computer network attack and defense tools takes another highly-specialized skill set that is not commonly found in university faculty.

To facilitate the creation of modules, as well as using the PEN as a platform to demonstrate these modules, we believe the right path is a framework in which course demonstrations and lab exercises can be performed in a manner that is repeatable, reversible, extensible, and as automated as possible for ease of use. With such a framework in place, energy can be focused on doing the detailed research into new malware artifacts to produce modules, such as the SSH exploit described in Section II.

A. *Repeatability*

Demonstrating how attack programs exploit vulnerable services, but do nothing against hardened services, requires three basic elements: (1) a vulnerable service running on its required operating system; (2) an exploit program that takes advantage of the vulnerable service to obtain unauthorized access, or elevate privileges, on the target host; and (3) a patch that can be applied to render the same service invulnerable to exploitation.

B. *Reversibility*

Using virtual machines, one can pre-configure the operating system and service, and can even mark the virtual machine as "immutable," permitting temporary modification of the system at run-time to demonstrate the effect of patching, while also preventing permanent changes to the virtual machine's file system. Alternately, one can simply copy bit-images of hard drives to restore pre-configured operating system/application sets. This establishes the "clean" state for attack and target platforms. This is, relatively speaking, the easy part.

From this point on, any activity on either the attack or target platform creates changes to the file system: new files, deleted files, new log entries (or destruction of system logs altogether by the attack tools), etc. Some attack software may entirely destroy the target host, rendering it impossible to then patch the vulnerable service at all.

In many cases, students benefit most from meaningful repetition in a safe environment where they can have a hands-on learning experience. This calls

ISBN 1-933510-96-7 / $15.00 ♡ 2008 CISSE

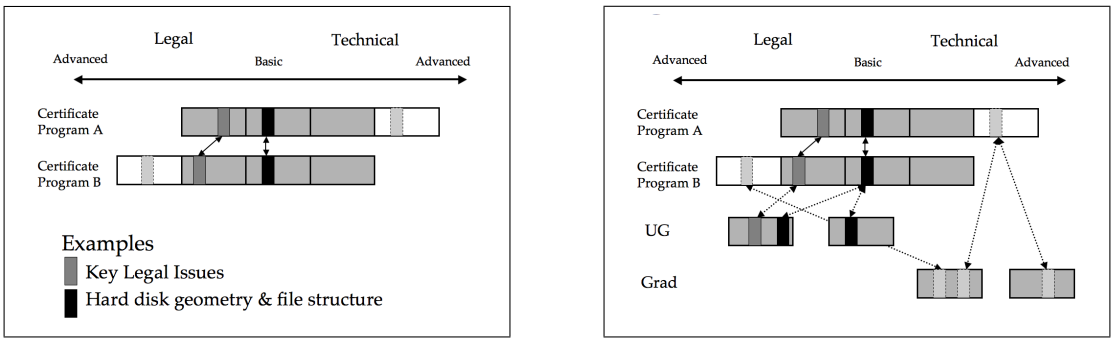

(a) Common modules within separate programs

(b) Re-use of modules shared by programs

Figure 4: Effect of module re-use in other programs

for a simple way to "turn back the clock" – reverting to an earlier state on demand – in order to predictably repeat the experiment. The more work involved in setting up the demonstration, the longer it takes and greater the chance for random variation between runs.

C. *Extensibility and ease-of-use considerations*

Development of a demonstration, as described in earlier sections, can be very time consuming and requires expertise. In order to develop a sufficiently broad set of modules, to maintain them over time, and to add new modules as new exploits or new security tools are developed, the framework must support modularity of design and be easily extended. This can be done using a framework involving a combination of configuration files and command line scripts that automatically organizes and presents faculty with a menu of options from which to chose, including automating and documenting the use of the system and PEN hardware itself. Basic techniques of automated system administration and software development will be used, so new modules can be easily created, tested, distributed, and added to the framework.

IV. Implementation Challenges

There are several challenges that must be met in order to attain the goals of creating a 10 year learning path as described thus far. We will consider them in this section.

A. *Risk aversion*

One factor slowing the adoption of novel educational programs is risk aversion, or the tendency to pursue only those activities already proven to bring predictable financial rewards. For example, it is very common to see standard computer forensics certificate programs at a large percentage of the NSA Centers of Excellence. There is very little difference between these programs. At the same time, there are almost no programs that focus on advanced incident response and reverse engineering of malware. Any institution attempting to create such a program runs into the problem of trying to branch out into a new area possessing no proven track record, no guarantee of students, and no clear job descriptions in government or business. Yet the lack of highly skilled incident response professionals is part of the reason that there are so many successful and long-term security incidents today, and an acknowledged lack of capacity to respond to a potential cyberwarfare campaign. [16]

When trying to convince university faculty and administrators to embark on creating new and creative learning opportunities, a question that often arises is "where is the market for students who graduate from this new program?"

The number and frequency of security incidents seen today supports a conclusion that there is not enough capacity today to adequately secure computer networks. This points to there being a potential demand, however the demand is not realized typically until *after a breach occurs*, at which point it is too late. A reasonable conclusion can be drawn that a market for highly trained security professionals does not fully exist, due in part to the level of pain from breaches not yet being fully felt by executive level management on down through the management ranks. It may take a catastrophic event which cannot be swept under the carpet, is not covered by insurance,

ISBN 1-933510-96-7 / $15.00 ♡ 2008 CISSE

or so shakes the customer base that the existence of the corporation as an entity is threatened, in order for a change in hiring and management practices to take place that increases the demand for skilled security professionals.

Once enterprises are compromised, and begin to react by reorganizing to include IA-aware staff, a classic *catch-22'*situation follows: a new graduate cannot get a job without first gaining enough work experience, yet cannot gain work experience without first getting a job.

B. Attaining sufficient real-world work experience

A significant hurdle within today's academic environment is the limitation on time available to students for gaining practical experience over general conceptual or theoretical learning. At best, undergraduates typically can only get about 1-2 years of practical experience, and often only one academic quarter of that is focused in depth.

For example, capstone projects often bring small teams of students together to develop prototype software applications or business plans for a novel concept. These are typically one academic quarter in length, during the senior year. Practical experience at the graduate level is often more in-depth, with year-long research projects in the final year of a Masters Degree program being common. These programs may be sponsored by businesses, who bring real-world problems to the table, and fostering solutions developed by the graduate student development team that may end up being used in production settings.

What is not common today are educational or research efforts that span undergraduate to graduate programs, or more elaborate programs that bring business or government and academia together on a multi-year basis. For example, it may be possible to get 5-6 years of experience with a voluntary service program tied to educational funding, or a form of *scholarship for service* while the student is still in school [17], provided the students start in their undergraduate studies and continue on immediately to higher degree programs. Such a program might work as follows.

Because there is insufficient computer security incident response capacity in most schools in the U.S. today, there is an opportunity to solve three problems simultaneously. First, a pool of students who are studying CNO can have an opportunity to apply those skills to proactively securing computers on their school's network, and to assist when they come under live attack. This is real-world experience dealing with actual attacks. Second, the staff and faculty who operate those computers get additional assistance in proactive and reactive computer security services without having to hire dedicated staff for this purpose.[3] Lastly, if there was scholarship support and a small stipend accompanying this effort, students would could provide these proactive and reactive security services instead of seeking other non-IA related employment opportunities simply to pay for their schooling.

Everyone would win in such a scenario, including the future employers of these students who would be able to hire new graduates who already possess several years of real-world experience by the time they complete their studies.

C. Type of skills being learned

A criticism often voiced by those hiring new graduates that they must immediately be re-trained to have the skills necessary to fit into the workplace. [18] They must either change the way the new-hire thinks to fit the reality of corporate needs, or to think in a different way than they were used to thinking while in school. This is equally true for software engineering as it is for CNO.

This may call for a shift in focus to more applied research, with less of a focus on theory and develop of new avenues of research. Or it may call for an increase in the use of internships and externships that are closely tied to the educational curriculum so as not to have perceivable gaps between the classroom and the corporate environments. A voluntary service program, such as the one described earlier in this section, may yield the desired result. Another alternative is integrating a research laboratory or *think tank* that studies emerging threats with special topics courses in which students receive credit for studying these emerging threats as they are found *in the wild*.

V. Conclusions

In conclusion, the development of an under-graduate IA concentration with an eye towards integration with

[3]Research grants that fund the purchase of computer equipment typically do not include funding for operating system upgrades, system administration support capable of adequately securing and monitoring the systems for breach, or incident response support when these systems almost inevitably get compromised. In addition, university overhead consumes approximately 50% of grant funding for overhead expenses, which do not include incident response services.

ISBN 1-933510-96-7 / $15.00 ♡ 2008 CISSE

future advanced degree programs, including research at the masters and doctorate levels, is a realistic and attainable goal. It moves towards establishing a 10-year pathway towards attaining advanced CNO skills. Using a framework for rapid integration of new modules and lab exercises keeps the curriculum fresh and provides one avenue for maintaining a high level of challenge to students, while leveraging limited technical expertise to a large faculty population with non-IA backgrounds. If done in a creative way, by merging volunteer service benefitting the educational institution and local enterprises and/or government agencies, combined with scholarship support for the students, there can be a simultaneous benefits to all involved. Further integration with computer security research laboratories that study emerging CNO tools adds even more currency to the educational experience. An improvement in the security posture of the institutions involved can be achieved, as well as providing students with hands-on experience that is immediately applicable to their future career.

VI. Acknowledgements

The author would like to thank Richard Anderson and the anonymous reviewers for comments on, and input to, this paper.

VII. References

[1] Frontline. Cyberwar, 2003. http://www.pbs.org/wgbh/pages/frontline/shows/cyberwar/.

[2] Josh Rogin. Air Force leaders hold Cyber Summit, 2006. http://www.fcw.com/article96881-11-17-06-Web.

[3] Office of Management and Budget. Implementation of Commonly Accepted Security Configurations for Windows Operating Systems, 2007. http://www.whitehouse.gov/omb/memoranda/fy2007/m07-11.pdf.

[4] Ross Kerber. TJX agrees to reimburse banks, 2007. http://www.boston.com/business/globe/articles/2007/12/01/tjx_agrees_to_reimburse_banks/.

[5] Linda Dailey Paulson. Wanted: more network-security graduates and research, 2002. http://ieeexplore.ieee.org/iel5/2/21180/00982909.pdf.

[6] Philip E. Ross. The Expert Mind, 2006. http://scientificamerican.com/print_version.cfm?articleID=00010347-101C-14C1-8F9E83414B7F4945.

[7] Warren Olney. The US and Unintended Consequences in Iraq, 2008. http://www.kcrw.com/news/programs/tp/tp080114the_us_and_unintende.

[8] Executive Office of the President. Presidential Decision Directive 63, 1998. http://www.ciao.gov/resource/paper598.pdf.

[9] Cynthia E. Irvine, Shiu-Kai Chin, and Deborah Frincke. Integrating Security into the Curriculum, 1998. http://www.cs.nps.navy.mil/people/faculty/irvine/Publications/Publications1998/IntegratSecCurric_Computer98.pdf.

[10] Melissa Jane Dark, Joseph J. Ekstrom, and Barry M. Lunt. Integrating Information Assurance and Security into IT Education: A Look at the Model Curriculum and Emerging Practice, 2006. http://jite.org/documents/Vol5/v5p389-403Dark127.pdf.

[11] Milton Greenberg. A University Is Not a Business (and Other Fantasies). *EDUCAUSE*, 2004. www.educause.edu/ir/library/pdf/erm0420.pdf.

[12] Tim Rosenberg and Lance Hoffman. Taking the network on the road, 2004. http://www.cpi.seas.gwu.edu/library/docs/2004-05.pdf.

[13] David Dittrich. Analysis of SSH CRC32 compensation attack detector exploit, 2001. Available at http://staff.washington.edu/dittrich/misc/ssh-analysis.txt.

[14] Michael Whitman and Herbert Mattford. A Draft Model Curriculum for Programs of Study in Information Security and Assurance. In *Proceedings of the 8th Colloquim for Information Systems Security Education*, 2004. http://www.ncisse.org/publications/cissecd/Papers/S3P04.pdf.

[15] Colin Armstrong and Nimal Jayaranta. Teaching computer forensics: Uniting practice with intellect. In *Proceedings of the 8th Colloquim for Information Systems Security Education*, 2004. http://www.ncisse.org/publications/cissecd/Papers/S4P03.pdf.

[16] Bob Brewin. U.S. unprepared for ongoing cyberwar, say top military and intelligence officials, March 2008. http://www.govexec.com/story_page.cfm?articleid=39466&dcn=todaysnews.

[17] Barack Obama's Plan for Universal Voluntary Public Service, 2008. http://www.barackobama.com/issues/service/.

[18] W. S. Curran. Teaching software engineering in the computer science curriculum. *SIGCSE Bull.*, 35(4):72–75, 2003.

ISBN 1-933510-96-7 / $15.00 ♡ 2008 CISSE