Journal of Information Technology Education: Research **Volume 12, 2013**

Cite as: Pittman, J. (2013). Understanding system utilization as a limitation associated with cybersecurity laboratories – A literature analysis. *Journal of Information Technology Education: Research, 12,* 363-378. Retrieved from http://www.jite.org/documents/Vol12/JITEv12ResearchP363-378Pittman0440.pdf

Understanding System Utilization as a Limitation Associated with Cybersecurity Laboratories – A Literature Analysis

Jason Pittman
Information Sciences, Capitol College, Laurel, Maryland, USA

jmpittman@capitol-college.edu

Executive Summary

The use of laboratories as part of cybersecurity education is well evidenced in the existing literature. We are informed about the benefits, different types of laboratories and, in addition, underlying learning theories therein. Existing research also demonstrates that the success of employing cybersecurity laboratory exercises relies upon controlling or minimizing potential challenges in the maintenance and usage of these laboratories. However, to date there has not been any effort to examine the possible limiting factors associated with system utilization in such laboratories.

In order to begin addressing such laboratory system resource utilization factors, this study examined existing research with the objective to gain an understanding of system utilization as a limitation to employing cybersecurity laboratories as pedagogical tools. Understanding the potential issues and limiting factors is of benefit to researchers, laboratory designers, and managers, as well as to higher education institutions hosting such laboratories. To this end, this study analyzed 11 years of academic literature for themes of limiting factors related to system utilization.

Analysis of the academic literature detected a reoccurring presence of system utilization issues as limiting factors within both hardware-based laboratories as well as virtualized laboratories. Furthermore, findings indicated that despite various attempts to resolve such utilization issues, the literature has yet to definitively do so. Concurrently, this study developed a taxonomy to conceptualize the evolution of, and relationships between, such literature.

The full taxonomy is constructed over the course of seven diagrams and demonstrates the evolution of both the utilization issues documented in the literature as well as potential solutions. The diagrams incorporate successive groupings of literature into six tiers which are arranged according to the emergent themes. According to each emergent theme (e.g., transitioning from hardware-based laboratories to virtualized laboratories and associated, residual utilization issues), connections within the tier as well as external to the tier (forwards and backwards) are established.

Ultimately, despite the success and effectiveness of cybersecurity laboratories, challenges related to system resource utilization continue to plague educators. What is more, no singular laboratory infrastructure design presents a model to address the system resource challenges. Thus, a need exists to create a cybersecurity laboratory design that

Editor: Benson Soong

minimizes system resource utilization limitations now. Such designs may be all-encompassing or potentially address only a limited collection of cybersecurity scenarios.

Keywords: Cybersecurity laboratory, laboratories, system utilization, literature analysis

Introduction

Colleges and universities are facing increased demand to produce cybersecurity graduates with real word knowledge (Yang et al., 2004). According to Locasto, Ghosh, Jajodia, and Stavrou (2011), the demand for skilled cybersecurity professionals is coming from both industry and federal entities. Principally, in 2009, the President of the United States of America Barack Obama affirmed the need to expand the nation's cybersecurity workforce ("The Comprehensive National Cybersecurity Initiative"). In response, government entities have begun fervently recruiting and training cybersecurity professionals. Moore (2011) reported that the National Security Agency (NSA) would seek as many as 3,000 new cybersecurity professionals. Additionally, the Homeland Security Department requested in the Financial Year 2013 appropriations for $12.9 million alone for cybersecurity education (O'Connell, 2011).

Academics also recognize the growing demand for skilled and capable security professionals (Bratus, Shubina, & Locasto, 2010; Maconachy & Gibbs, 2009; Mirkovic & Benzel, 2012). There exists tacit acknowledgement of a lack of individuals graduating from institutions of higher education with practical security knowledge. Locasto et al. (2011) in particular lamented that not enough trained security professionals are graduating from higher education. Likewise, Harmon (2011) indicated that the need for knowledgeable security professionals would continue into the near future. Such commentary from academia aligns with the overarching call from both the president as well as the federal workplace.

One prominent method academia uses to provide practical, hands-on security education is laboratory exercises (Duffany & Cruz, 2012; Mattord & Whitman, 2004). Anantapadmanabhan, Memon, Frankl, and Naumovich, (2003) stated that cybersecurity laboratory exercises are a prominent means for learners to acquire a meaningful cybersecurity higher education. Indeed, research by McKinney (2010) as well as Kaucher and Saunders (2002) supports the notion that the laboratory exercise based approach to learning is superior to traditional methods alone (e.g., lecture). Yet, employing laboratory exercises is not a trivial endeavor.

The success of employing cybersecurity laboratory exercises to provide practical experience to higher education learners relies upon controlling or minimizing potential challenges. According to Stockman (2003) one such challenge of particular concern is laboratory system utilization (e.g., CPU, memory, and disk utilization). The concern stems from conflict between the importance of laboratory exercises in the acquisition of practical cybersecurity knowledge (Bhagyavati, 2006) and the potential for system resource issues to directly limit the number of learners able to access these laboratory environments (Stewart, Humphries, & Andel, 2009). Unfortunately, there is little or no overarching review of the literature discussing such limitations associated with cybersecurity laboratories.

In order to begin addressing such laboratory system resource utilization issues, a survey of the exiting literature may be of benefit. Thus, the intent of this study was to gain an understanding of existing research related to system utilization as a limitation to employing cybersecurity laboratories as pedagogical tools. Through analysis of existing literature, this study may provide cybersecurity laboratory developers and implementers with a basis from which improvements in cybersecurity laboratory system utilization can be achieved. As well, a topology of the relationships between the included studies may be of benefit to other researchers exploring the evolution of cybersecurity laboratories. The topology is represented in a series of graphical figures and separated

into *tier* by color. Furthermore, examination of the literature may yield recommendations for future research related to techniques or technologies that reduce system utilization.

Title Searches, Articles, Research Documents, and Journals

This study included a search of available literature addressing laboratory design in cybersecurity education with a focus on system resource issues and possibly remediation. Primary sources of related literature included ProQuest and Academic Search Premier (EBSCOHost). This study performed a search of the ProQuest dissertation and these databases in order to affirm that the study has not previously been undertaken. Additionally, literature searches through Google™ Scholar and Google™ Books provided access to Internet journals, papers, conference proceedings, and books. Where online sources failed to provide accurate citation information, publishers or authors were contacted in order to obtain missing citation information.

The literature search used keywords, key phrases, and search operators during the literature search. Examples of the keywords and key phrases included, but were not limited to *information assurance*, *information security*, *cybersecurity, laboratory, lab, exercises, hands-on, key performance indicators, utilization, system resources, challenges,* and *recommendations*. Variations of key words were utilized to ensure due diligence and to focus results on research that discussed the system utilization issues associated with cybersecurity laboratories. Moreover, the earliest research discovered during the literature review was Clark (2001). Therefore, the search date range covered the 11 years in order to give full treatment to the body of literature under examination. The year 2013 was excluded from the search because the year is not yet completed. The full date range was divided into two periods for the purpose of analysis: the background of cybersecurity laboratory system utilization issues (from 2001 to 2006) and more recent cybersecurity laboratory research (from 2007 to 2012).

Literature Topology and Visualization

Literature was grouped into tiers according to *research theme* (e.g., hardware-based laboratories) in order to facilitate easier conceptualization of the research and to enhance understanding of the progress of system utilization limitations over the course of the literature.

The literature topology consists of a series of circles that are separated into tiers. Tiers are defined by color groupings; circles of the same color, oriented roughly in the same vertical, belong to a single tier of literature. Tiers are a means of categorization relative to the surrounding literature: a means to visually distinguish between otherwise textual themes emerging during the literature analysis. Circles in the same tier collectively represent a shift in the literature both forward in time as well as forward in attempts to address system utilization limitations.

Size of individual circles imparts no meaning in the topology. Connections between individual studies in the topology do impart meaning. The arrow on the line indicates which study extends the other (without the arrow). For example, if *B* extends A the relationship would be seen as A → B. Connections appear in three forms: full line connection, dotted line connection, and no connection. The single exception to these three forms is the first circle in the first tier. The first circle in the germinal tier represents the seminal study and displays all connections, potential and realized.

Those studies appearing with full line connections *strongly* extend the existing research. Strength in this context was observed as the appearance of extensive reference (citation) and through a detailed level of related subject material (e.g., CPU utilization in hardware-based laboratories). Studies connected through a dotted line share a *weak* relationship. Weak relationships were observed as the appearance of limited reference to existing literature or through a conceptual level of related subject material (e.g., cybersecurity laboratory exercises). Literature appearing with no

connections tangentially extends the surrounding research. The tangential nature of the contribution to the field was observed as having very limited or no reference to existing literature and starkly different solution to the system utilization problem.

Background of Cybersecurity Laboratory System Utilization Issues

The background of system utilization issues in cybersecurity laboratories began with Clark (2001). The significant contribution from Clark comes from the extensive discourse of laboratory infrastructure design and implementation. In particular, Clark commented that laboratory infrastructure should include the fastest CPUs, fastest and largest hard disks, and largest amount of memory possible. Such represented a tacit understanding of the need for adequate system resources (i.e., CPU, memory, and disk utilization) in order to support the laboratory exercises. By proxy, such also establishes the fundamental points of limitation for laboratories. Accordingly, Figure 1 demonstrates *graphically* Clark's research as the foundation of literature associated with limitations in cybersecurity laboratories. Finally, Clark also made brief reference to an early version of standalone VMware but (a) did not associate any potential issue resolution to the use of virtualization and (b) concluded that multiple hardware systems were preferred due to the cost of the VMware solution.

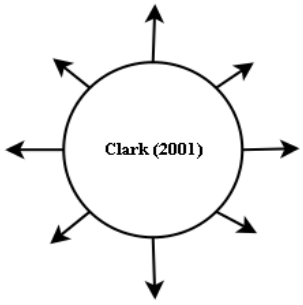

***Figure 1*. This figure illustrates the seminal study concerning limitations of cybersecurity laboratories. Connections represent connections to all potential and realized literature.**

First tier research

Several studies followed closely after Clark (2001) and further discussed limitations of cybersecurity laboratories. These studies exhibited a direct relationship to Clark either through reference or context. Figure 2 documents the direct relationship in the research between Clark (2001), Schafer, Ragsdale, Surdu, and Carver (2001), Hill, Carver, Humphries, and Pooch (2001), as well as Micco and Rossman (2002). As well, this tier of research focused on hardware-base cybersecurity laboratories and attempted to resolve utilization issues by sizing the technical infrastructure up or down.

Schafer et al. (2001) identified a need akin to Hill et al. (2001) and responded by designing and implementing a cybersecurity laboratory analogous in purpose to “conventional weapons training” (p. 226). The concept of isolation, taken from conventional weapon firing ranges, addresses the shortcoming of shared laboratory spaces such as Hill et al. (2001) discussed. The laboratory infrastructure Schafer et al. (2001) implemented was composed entirely of hardware based systems on the scale of an enterprise environment. To this effect, the authors commented that a smaller, less complex laboratory setup would have been sufficient. The size of this smaller laboratory environment corresponded to Hill et al. (2001) in that the total number of systems will be 5 to 8 systems.

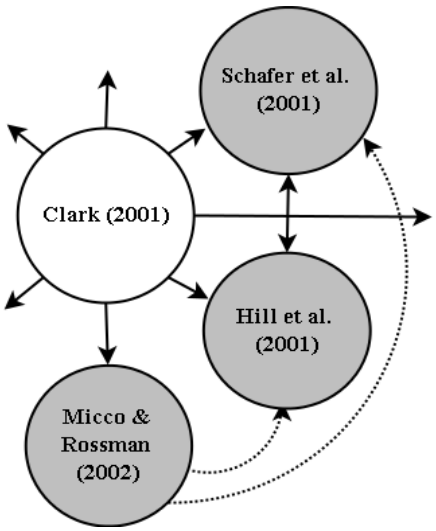

***Figure 2*. The first tier of research extending the seminal work of Clark and focusing on utilization limitations in hardware-based laboratory environments.**

Similar to Clark (2001), Micco and Rossman (2002) considered virtualizing laboratory cybersecurity infrastructure with VMware. Micco and Rossman determined that virtualization would add unnecessary complexity and, like Clark (2001), elected to use multiple hardware based systems. Still, Micco and Rossman constructed a cybersecurity laboratory that according to the authors properly emulated a realistic environment. Attempting to provide such an authentic experience before developing the laboratory exercises represents a departure from previous research. The authors described in detail the 7 systems that comprised the infrastructure (Micco & Rossman, 2002). This number closely mimicked the design displayed in Hill et al. (2001) and Schafer et al. (2001). Moreover, Micco and Rossman (2002) demonstrate an early understanding of what the minimum laboratory infrastructure makeup must be to present a realistic architecture to learners. Another interesting aspect to Micco and Rossman was the exposure to the obstacles, mistakes, and challenges encountered during the implementation of the cybersecurity laboratory. One possibility is that due to the maturation of cybersecurity laboratory research in general, sharing experiences and lessons learned represents an impending shift in the literature.

Second tier research

The second tier of research, shown in Figure 3, represented a partial move away from hardware-based cybersecurity laboratories and into virtualization technology. However, the trepidation of moving into fully virtualized laboratories was apparent as much of this tier of research focused on identifying the differences between hardware-based and virtualized laboratories as well as isolated use of virtualization. While one study recognized the physical limit of hardware-based laboratories with respect to space, several other studies in this category discussed potential virtualization utilization issues exclusively. Still, the literature in the second tier proved to be instrumental as the basis for eventual transition to fully virtualized laboratories.

Pushing forward on the design front, Padman and Memon (2002) broke new ground in two important aspects. One such aspect was that the authors discussed a general issue present in existing cybersecurity laboratory design and architecture. The other aspect in which Padman and Memon broke new ground was in employing virtualization to address observed shortcomings in those hardware based laboratories. Specifically, the authors asserted that cybersecurity laboratories are "difficult to build and maintain" (p. 1). In order to create context for such innovation, Padman and Memon outline a series of core characteristics that cybersecurity laboratories should exhibit. Chief amongst these characteristics are scalability, cost effectiveness (in particular, maintenance),

robustness, realism, and finally maintainability. The significance of Padman and Memon expressing these characteristics is twofold: (a) such builds upon prior research (Clark, 2001; Schafer et al, 2001; Hill et al., 2001) and (b) establishes discrete concepts for identifying shortcomings in cybersecurity laboratories. Collectively, these advances in the literature laid the foundation for a transition between historical research germinal to the cybersecurity laboratory domain and current research.

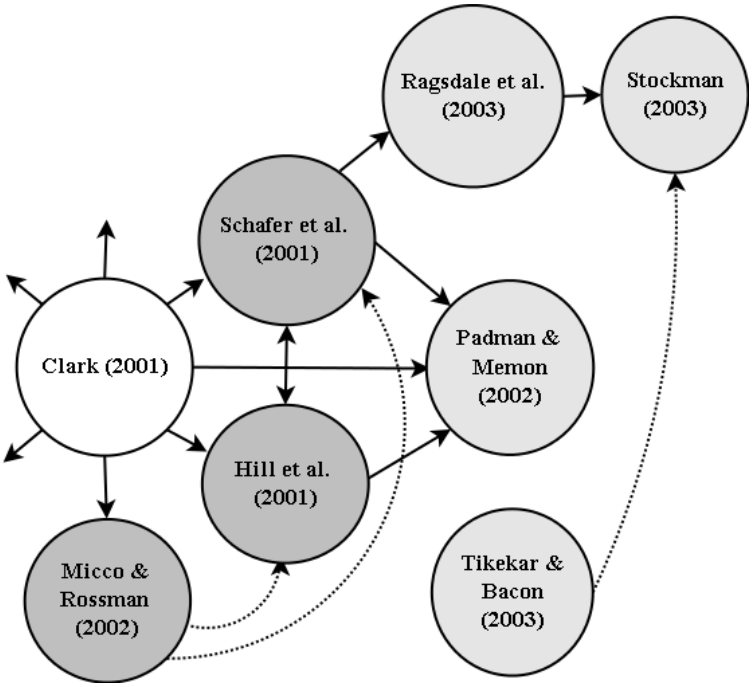

***Figure 3*. The second tier of research beginning to bridge from hardware-based laboratories to virtualization as an attempt to address utilization limitations.**

Following Padman and Memon (2002), Tikekar and Bacon (2003) indicated that challenges exist in the creation and implementation of cybersecurity laboratory exercises. Such challenges included learners needing to share laboratory systems due to class sizes outpacing laboratory availability, the maintenance of the laboratory systems (e.g. setup, configuration, and restoring baselines after exercise completion,), as well as laboratory exercises requiring more time to complete. These challenges are likely interrelated insofar as limited hardware resources directly bound the number of learners that can use the laboratory at any one time and bound the time necessary to complete CPU, memory, and disk intensive exercises.

Ragsdale, Lathrop, and Dodge (2003) described additional research based on Schafer et al. (2001). The chief additions Ragsdale et al. (2003) offered were around the technological infrastructure of the laboratory. Foremost, Ragsdale et al. described the migration of a laboratory infrastructure to a virtualized design. The authors claimed that non-virtualized laboratory environments demand "significant investments in terms of hardware, software, and human resources to build and maintain the physical networks of computers and communication components" (p. 3). Additionally, Ragsdale et al. asserted that utilizing multiple hardware based systems could be physically untenable in some laboratory infrastructure designs. The validity of such an assertion was relative to both the scale of the laboratory (in terms of the number of systems presented) but is also relative to the number of educators available to maintain and manage a laboratory of any size. In addition, such issues are similar to the shortcomings identified by Padman and Memon (2002).

Stockman (2003) continued the example of Ragsdale et al. (2003) and created a cybersecurity laboratory that leveraged virtualization technology (e.g., VMware). Apart from the use of virtualization, Stockman described a standard laboratory approach as previously demonstrated in research such as Clark (2001), Micco and Rossman (2002), and Schafer et al. (2001). However, Stockman provided keen insight into a challenge that future research would face, namely, that virtualized laboratory infrastructure has an inherent limitation on the number of concurrent virtual systems that may run based on the virtualization host's CPU, memory and disk utilization. Although Stockman did not make this statement in the context of addressing a problem or gap in the existing literature, the assertion does echo comments from Clark (2001) and Tikekar and Bacon (2003), thus establishing an early corollary between common issues and limitations between hardware based and virtualized cybersecurity laboratories.

After the studies published in 2003, the nature of cybersecurity laboratory research began developing along several new angles. Bishop and Frincke (2004) noted that "cybersecurity education has arguably left the pioneer stage" (p. 63). Based on the literature review, it appeared that researchers were becoming more aware of the challenges and limitations associated with the hardware based approach. At the same time, technology continued to evolve and virtualization developed into a stable, mature platform. Thus, as challenges and limitations arose in hardware based laboratory designs researchers increasingly turned to virtualization. However, the literature review also revealed that virtualization presented some of the challenges and along with introducing some new limitations.

Third tier of research

Figure 4 shows the small (three studies) but impactful expansion of the literature in the third tier of research. The impact of this tier of research was evident primarily in how the studies herein directly extended the prior tier. Moreover, one of the studies was directly support by the other two in the same tier. All three of the studies examined utilization issues with fully virtualized cybersecurity laboratories, thus preparing the way for a later fourth tier to attempt to resolve such issues.

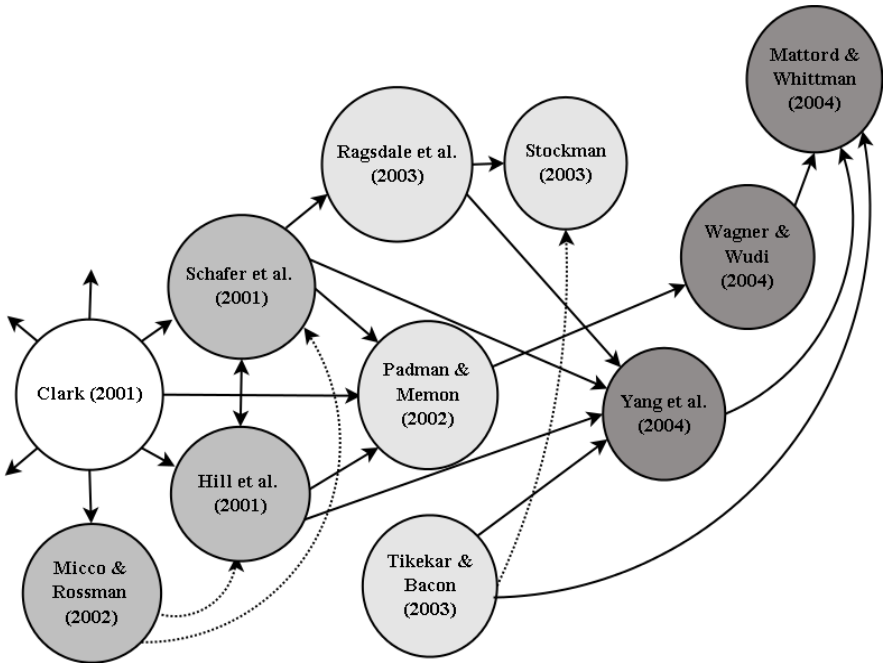

Figure 4. **The third tier of research explores virtualization more in depth as a potential solution to utilization issues.**

Yang et al. (2004) summarized the cybersecurity laboratory research of such foundational literature as Hill et al. (2001), Schafer et al. (2001), and Tikekar and Bacon (2003). The summary however did not take the form of a literature review. Rather, Yang et al. (2004) used the previous literature to establish 8 categories of challenges or shortcomings. Five of these categories were related to potential issues such as maintaining cybersecurity laboratories, system resources (both in the context of laboratory systems being able to run the laboratory exercises as well as the laboratory having enough systems for the learner population), and the ability to accurately mimic real world infrastructure. The remaining two categories included learner access to cybersecurity laboratory facilities and access to the Internet from within laboratories. The solution proposed by Yang et al. was elegant and closely resembled the laboratory infrastructure outlined by Schafer et al. (2001). Accordingly, Yang et al. (2004) opted to use more hardware based laboratory systems distributed across a large physical space. Doing so ignored the issue identified in Ragsdale et al. (2003) regarding the physical space limitations associated with hardware based laboratory designs.

The research of Wagner and Wudi (2004) centered on the challenges of managing a cybersecurity laboratory and the laboratory's system resources. While the majority of challenges were related, learner access to laboratory systems and learners' lacking foundational knowledge, the authors made a clear point of the potential benefits of virtualization. Specifically, Wagner and Wudi explained that using a virtualization solution such as VMware would allow cybersecurity laboratories to (a) support additional laboratory systems and (b) assist educators in more easily managing the laboratory systems. The discussion of using virtualization to ease or eliminate management burdens echoed prior work by Padman and Memon (2002). A final observation is that the laboratory infrastructure Wagner and Wudi (2004) utilized started with 8 hardware based systems. Approximations of this number appeared several times in the foundational literature as well (Hill et al., 2001; Schafer et al., 2001; Micco & Rossman, 2002).

Mattord and Whitman (2004) stated that educators must "understand the unique demands" of designing and implementing a cybersecurity laboratory (p. 8). Mattord and Whitman present these challenges in a series of best practice recommendations. The intent of the authors was to provide educators and researchers with a number of possible choices when designing and implementing a cybersecurity laboratory. The best practices extended the arguments set forth in Wagner and Wudi (2004). Similar to Wagner and Wudi, along with Yang et al (2004), Mattord and Whitman (2004) outlined a number of issues related to laboratory access, operating system selection, and laboratory network segmentation. However, where Mattord and Whitman contributed a new perspective to the body of research is in hardware recommendations and virtualization options. Mattord and Whitman discussed the need of "highly capable" hardware based servers (p. 9). Although the motivation for such recommendations was not explicit in the study, it is a reasonable assumption to conclude that Mattord and Whitman recognized the potential issues associated with insufficient system resources (e.g., CPU, memory, and disk utilization). The rationale for this assumption rested in Mattord and Whitman citing prior work where others did make those issues explicit (Tikekar & Bacon, 2003; Wagner & Wudi, 2004). Furthermore, when discussing virtualization, Mattord and Whitman (2004) indicated two challenges: memory utilization as a limitation and, like the hardware concept, the need for robust servers.

Fourth tier of research

In the fourth tier of research, the literature continued to expand on issues in cybersecurity laboratory virtualization and potential methods of addressing such issues. As a departure from previous tiers, the fourth tier focused on mapping utilization issues associated with various hardware-based, partial virtualization, and full virtualization solutions. The fourth tier, as seen in Figure 5, also demonstrated early use of virtualization hosts to serve multiple, concurrent laboratories.

However, the virtualization host technique introduced new utilization issues which Bullers, Burd, and Seazzu (2006) cataloged.

Villanueva and Cook (2005) gave further explanation for the challenges inherent in virtualization. Villanueva and Cook analyzed three broad options affiliated with delivering a cybersecurity laboratory implementation. These options consisted of using a hardware based infrastructure, multiple virtualized systems, and a virtualization server host. The authors considered the hardware based approach to be untenable due to the requisite labor of upkeep and physical space considerations. The authors decided to use a virtualization server host based on the increased capability to manage laboratory systems. The increase in management ease came from the centralized access to individual laboratory systems through the virtualization host. Notwithstanding the benefits, Villanueva and Cook did note that the virtualization host option incurred high financial cost in addition to a different set of maintenance concerns. More importantly, Villanueva and Cook detailed the CPU, memory, and disk limitations of using the virtualization host solution. The authors explained that laboratory exercises had to be restricted to a small number of virtual machines due to high CPU, memory, and disk utilization.

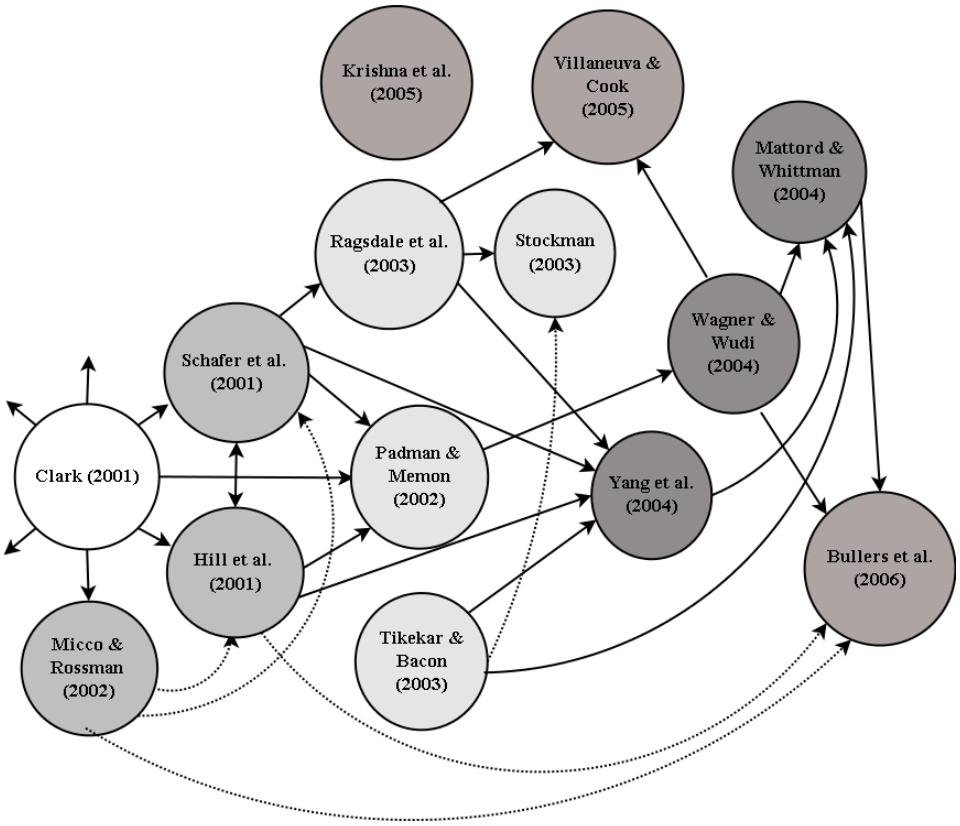

Figure 5. **The fourth tier of research representing a more complete move to virtualized laboratories and a robust mapping of utilization limitations to potential solutions.**

Like Yang et al. (2004), Krishna, Sun, Rana, Li, and Sekar (2005) identified a series of general issues with developing a cybersecurity laboratory. These issues collectively fall under the auspices of maintenance and management challenges (Ragsdale et al., 2003; Villanueva & Cook, 2005; Yang et al., 2004). Initially, Krishna et al. (2005) sought to leverage a virtualization host server Akin to Villanueva and Cook (2005). In addition, as did Villanueva and Cook (2005), Krishna et al. (2005) observed that virtualization host CPU, memory, and disk utilization were bounding factors. Ultimately, Krishna et al. opted to use hardware based systems, each running a virtualization host that provided learners with approximately 6 to 9 virtual systems.

The shift from hardware based cybersecurity laboratories to virtualized infrastructures culminated in Bullers et al. (2006). Bullers et al. began by analyzing prior significant research in both hardware based infrastructure and virtualized infrastructure (Hill et al, 2001; Mattord & Whitman, 2004; Micco & Rossman, 2002; Wagner & Wudi, 2004). Bullers et al. (2006) concluded that realistic hardware based cybersecurity laboratories were no longer viable without outside funding due to the scale of designs and related upkeep. With respect to the virtualized infrastructure approach, Bullers et al. balanced the advantages and disadvantages of adopting the new type of laboratory architecture. The advantages grouped around alleviating upkeep and management. According to Bullers et al., the shortcomings of the virtualization approach were apparent in the networking support and disk space requirements. Whereas others such as Mattord and Whitman (2004) and Villanueva and Cook (2004) encountered challenges in the virtualization approach connected to CPU, memory, and disk utilization, Bullers et al. (2006) did not report any such findings. Based on the laboratory exercises Bullers et al. described, it is possible that the authors did not confront such system resource issues as a result of careful consideration and planning before implementing the laboratory infrastructure.

Modern Cybersecurity Laboratory Research

Stackpole (2008), representing the fifth tier of literature as seen in Figure 6, tangentially extended the work of Bullers et al. (2006) and provided an interesting discourse on the transition from hardware based laboratory to a virtualized design. What made the study particularly interesting was the amount of detail Stackpole exposed in assessing both laboratory infrastructure approaches. The hardware based concept, Stackpole claimed, will be too expensive in both financial

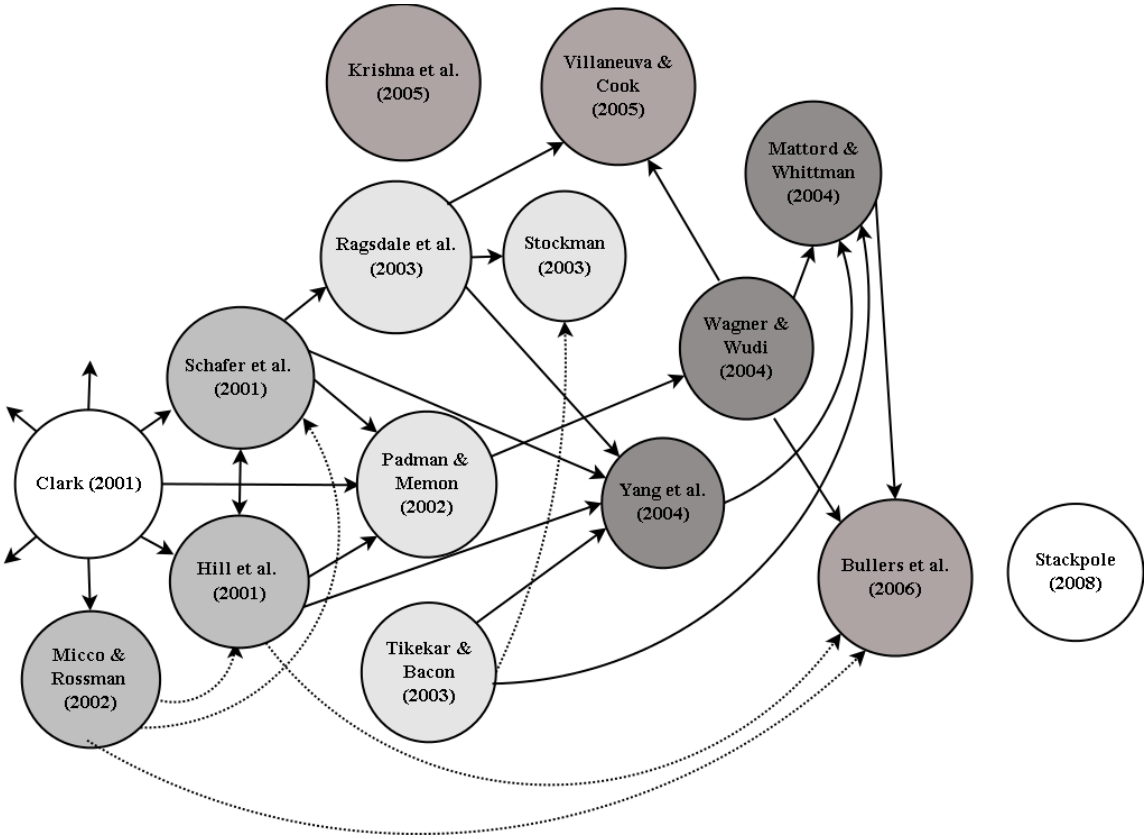

Figure 6. **The fifth tier as an entry point into the modern era of research in cybersecurity laboratory utilization issues represented by Stackpole (2008).**

terms and labor, unavailable too often, and have performance challenges. Stackpole conducted two pretest-posttest studies using a prototype laboratory infrastructure and volunteer learner participants. From the results, Stackpole concluded that a virtualized infrastructure using VMware resolved the two leading concerns from the analysis of the hardware based approach. However, according to Stackpole, the performance issues remained relative to CPU and memory utilization. Even more interesting, Stackpole suggested that additional, higher performance hardware might solve the CPU and memory utilization limitations. Weighing the evidence from the body of literature, such a suggestion does not seem likely to ultimately affect the CPU and memory utilization issues.

The sixth tier of research

The sixth tier, outlined in Figure 7, of research employed laboratory virtualization exclusively. Thus, the transition begun with Bullers et al. (2006) in the fourth tier and carried forward by Stackpole (2008) in the fifth tier was completed. Moreover, the studies began to address the utilization limitations present and perhaps residual from previous literature. A major distinguishing theme present across some research in this tier is the development of custom solutions external to the virtualization technology.

Conspicuously however, there are large gaps in the literature between the fifth and sixth tiers (and going forward). Not only are individual pieces of research farther apart temporally but there are demonstrably fewer studies in the modern era. As well, there is an overwhelming uniformity to the direction in which the sixth tier of research referenced other literature to a point where the sixth tier almost exists entirely separate from the previous tiers of literature.

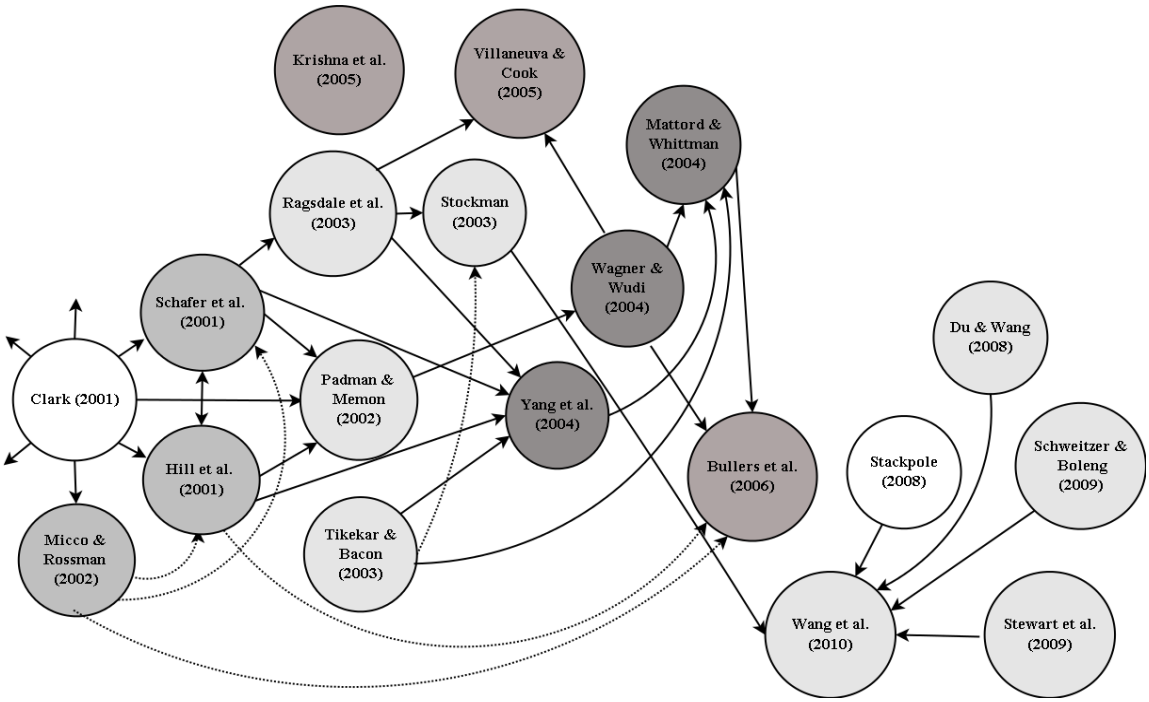

Figure 7. **The sixth tier of research representing a full conversion to virtualization with utilization issues still present and custom laboratory solutions.**

Complementary to the research focus of Du and Wang (2008), Schweitzer and Boleng (2009) investigated a cybersecurity laboratory approach not directly reliant on a hardware or virtualized infrastructure. Differing from Du and Wang (2008) however, Schweitzer and Boleng (2009) created a series of laboratory exercises through custom web applications. This approach negates the

shortcomings of hardware based and virtualization based laboratory infrastructures. Schweitzer and Boleng's web application laboratories closely resembled a combination of the tutorial and project methods of teaching cybersecurity espoused by Yurcik and Doss (2001). Schweitzer and Boleng (2009) argued that the strength of using web application laboratories rests in not requiring pre-existing knowledge of operating systems, tools, hardware, or virtualization packages. Du and Wang (2008) asserted a similar rationale, concluded similarly, but addressed the issue differently. However, whereas Du and Wang produced quantitative support, Schweitzer and Boleng (2009) offered anecdotal evidence for the effectiveness of web application laboratories.

Another study in the body of literature that contrasted and compared hardware based laboratory designs to virtualized approaches was Stewart et al. (2009). Stewart et al. discussed familiar pros and cons outlined in prior studies (Hill et al., 2001; Ragsdale et al., 2003; Wagner & Wudi, 2004). Where Stewart et al. enrich the body of literature is in adding emulation to the canon of laboratory types. Emulation, according to Stewart et al., was distinct from existing virtualized laboratory designs (e.g., VMware). The definition of emulation provided by Stewart et al. was remarkably similar to the honeypot technology described by Sadasivam, Samudrala, and Yang (2005). Given the attention Stewart et al. (2009) levy to the CPU, memory, and disk utilization challenges posed by both the hardware based and virtualization based laboratory approaches, a gap existed in not examining the potential of using a honeypot to emulate laboratory systems.

Wang, Hembroff, and Yedica (2010) extended previous research utilizing the virtualized infrastructure approach. Building on the preceding work (Du & Wang, 2008; Stackpole, 2008; Schweitzer & Boleng, 2009; Stewart et al., 2009; Stockman, 2003) the authors detailed how to implement VMware ESXi along with the Lab Manager add on package. The laboratory infrastructure – multiple virtualization hosts assembled in a cluster and comprised of large amounts of memory, fast CPU and disk storage – represented a more sophisticated design relative to previous research. According to Wang et al., Lab Manager made a significant impact towards resolving the management and maintenance of the virtual laboratory systems. At the same time, the authors reported CPU, memory, and disk utilization issues at times when learners were using the laboratory. Relative to the hardware supporting the ESXi environment, the limitations imposed by CPU, memory, and disk utilization is somewhat puzzling. However, the limitations correlate to prior studies (Krishna et al., 2005; Villanueva & Cook, 2005) and affirm the conjecture that such limitations exist in most types of cybersecurity laboratory implementations.

Conclusion

Laboratory exercises are an established and fundamental means of practical learning in cybersecurity education (Chatmon, Chi, & Davis, 2010; Du & Wang, 2008; Irvine, 1999; Yurcik & Doss, 2001). Within the body of research, hardware based laboratories and virtualized laboratories emerged as the dominant models of providing these exercises to learners (Bullers et al., 2006). The virtualization approach to laboratory implementation addressed much of the maintenance and access shortcomings of hardware based laboratories (Padman & Memon, 2002). However, later research (Tikekar & Bacon, 2003; Villanueva & Cook, 2005) revealed exacerbated issues in the virtualization approach related to CPU, memory, and disk utilization. Current research largely continued the virtualized laboratory implementation (Wang et al., 2010) while some innovation occurred that appears to escape the issues associated with virtualization (Du & Wang, 2008; Schweitzer & Boleng, 2009).

This study captured the evolution of cybersecurity laboratory design and associated challenges through an examination of historical and current studies. The effectiveness of laboratories in providing practical knowledge for cybersecurity learners was well evidenced (Chatmon et al., 2010; Du & Wang, 2008; Irvine, 1999; Yurcik & Doss, 2001). Such evidence came not only from learner feedback (Irvine, Warren, & Clark, 1997) but also from quantitative studies by Schafer et

al. (2001), Kaucher and Saunders (2002) as well as Du and Wang (2008). The balance of qualitative and statistical investigation provided a robust assessment of cybersecurity laboratory success.

Yet, despite the success and effectiveness of cybersecurity laboratories, challenges related to system resource utilization continues to plague educators. What is more, no laboratory infrastructure design alone appears to propose a model to address the system resource challenges. Resolving the system resource issues is particularly important when engaging learners in core cybersecurity exercises (Chatmon et al., 2010; Krishna et al., 2005; O'Leary, 2006). The inability to address the limitations resulting from utilization is compelling. In particular, the academic use of virtualization is not unlike the use of virtualization in IT-reliant industries. Therefore a natural inquiry to make is how do entities in various industries handle similar utilization limitations. Perhaps there is a different means of implementing existing technology, perhaps there is entirely new technology, or perhaps the answer is financially bounded for academia in a way that industry is not.

Regardless, a dire need exists to create a cybersecurity laboratory design that minimizes system resource utilization limitations now. Such designs may be all-encompassing or potentially address only a limited collection of cybersecurity scenarios. Existing technologies such as honeypots may be of potential benefit in reducing utilization. Likewise, potential designs need not be technological at all. There appears to be little or no discussion in the literature on the use of operations management to eliminate or reduce utilization limitations. If system resources are analogous to labor dollars, maximally efficient utilization might be reducible to an exercise of labor scheduling.

Lastly, there are the conspicuous gaps in the literature ranging from 2006 onward as discussed as part of the sixth tier of research. Based on reviews on complementary research- studies focused on cybersecurity laboratories but not necessarily *utilization issues* – such gaps may be related to the maturation of both the research field and the underlying technology. However, these gaps might also be related to a stagnation in innovation within the pedagogy or related to an acceptance of the utilization issues. Future research centered on understanding publishing trends, shifts in pedagogy, and student perceptions of laboratories in cybersecurity may be beneficial.

References

Ananthapadmanabhan, V., Frankl, P., Memon, N., & Naumovich, G. (2003, July). Design of a laboratory for information security education. In C. Irvine, & H. Armstrong (Eds.), *Security education and critical infrastructures* (pp. 61-73). Norwell, MA: Kluwer Academic Publishers.

Bhagyavati. (2006). Laboratory exercises in online information assurance courses. *Journal on Educational Resources in Computing, 6*(4), 1-5. doi:10.1145/1248453.1248457

Bishop, M, & Frincke, D. (2004). Joining the security education community. *IEEE Security & Privacy*, *2*(5) 61-63. doi:10.1109/MSP.2004.75

Bratus, S., Shubina, A., & Locasto, M. E. (2010). Teaching the principles of the hacker curriculum to undergraduates. In *Proceedings of the 41st ACM technical symposium on computer science education* (pp. 122-126). New York: ACM New York. doi:10.1145/1734263.1734303

Bullers, W. I. Jr., Burd, S., & Seazzu, A. F. (2006). Virtual machines – An idea whose time has returned: Application to network, security, and database courses. In *Proceedings of the 37th SIGCSE Technical Symposium on Computer Science Education* (pp. 102-106). New York: ACM Press. doi:10.1145/1124706.1121375

Chatmon, C., Chi, H., & Davis, W. (2010). Active learning approaches to teaching information assurance. In *InfoSecCD '10 2010 Information Security Curriculum Development Conference* (pp. 1-7). New York: ACM New York. doi:10.1145/1124706.1121375

Clark, P. C. (2001, May). Supporting the education of information assurance with a laboratory environment. In *Proceedings of the Fifth National Colloquium for Information Systems Security Education*, Fairfax, VA.

The Comprehensive National Cyber Security Initiative. (2009). *Executive Office of the President of the United States*. Retrieved from http://www.whitehouse.gov/sites/default/files/cybersecurity.pdf

Du, W., & Wang, R. (2008). SEED: A suite of instructional laboratories for computer security education. *Journal on Educational Resources in Computing, 8*(1), 1-24. doi:10.1145/1348713.1348716

Duffany, J. L., & Cruz, A. (2012). Design of a computer security teaching and research laboratory. In *Proceedings of the 43rd ACM technical symposium on Computer Science Education* (p. 678). New York: ACM New York. doi:10.1145/2157136.2157421

Harmon, D. (2011). *Careers in Internet security*. New York: Rosen Pub.

Hill, J. M. D., Carver, C. A. Jr., Humphries, J. W., Pooch, U. W. (2001). Using an isolated network laboratory to teach advanced networks and security. *Proceedings of the Thirty-second SIGCSE Technical Symposium on Computer Science* (pp. 36-40). New York: ACM New York. doi:10.1145/364447.364533

Irvine, C. E. (1999). Amplifying security education in the laboratory. In *Proceedings IFIP TC11 WC 11.8 First World Conference on Information Security Education* (pp 139–146). New York: Springer Publishing Company. Retrieved from http://cisr.nps.edu/downloads/papers/99paper_ampseced.pdf

Irvine, C. E., Warren, D. F., & Clark, P. C. (1997). The NPS CISR graduate program in INFOSEC education: Six years of experience. In *Proceedings of the 20th National Information Systems Security Conference* (pp. 22-30). Gaithersburg, MD: National Institute of Standards and Technology. Retrieved from http://csrc.nist.gov/nissc/1997/proceedings/022.pdf

Kaucher, C. E., & Saunders, J. H. (2002, June). *Building an information assurance laboratory for graduate-level education*, Paper presented at 6th National Colloquium for Information System Security Education, Redmond, WA.

Krishna, K., Sun, S., Rana, P., Li, T., & Sekar, R. (2005, June). *V-NetLab: A cost-effective platform to support course projects in computer security*. Paper presented at the 9th Colloquium for Information Systems Security Education, Atlanta, Georgia.

Locasto, M. E., Ghosh, A. K., Jajodia, S., & Stavrou, A. (2011). The ephemeral legion: Producing an expert cyber-security work force from thin air. *Communications of the ACM, 54*(1), 129-131. doi:10.1145/1866739.1866764

Maconachy, W. V., & Gibbs, M. G. (2009). Integrating cyber security into higher education curricula. In A. Aslant & F. Miah (Eds.), *Proceedings of the Third American Institute of Higher Education Conference 2(1)*, 513-518. Retrieved from http://www.amhighed.com/documents/Nashville2009/AmHighEd_Proceedings_Nashville_2009.pdf

Mattord, H. J., & Whitman, M. E. (2004). Planning, building and operating the information security and assurance laboratory. In *Proceedings of the 1st annual conference on Information Security Curriculum Development* (pp. 8-14). New York: ACM New York. doi:10.1145/1059524.1059527

McKinney, K. (2010). *Active learning*. Normal, IL: Center for Teaching, Learning & Technology.

Micco, M., & Rossman, H. (2002). Building a cyberwar lab: Lessons learned: Teaching cybersecurity principles to undergraduates. In *Proceedings of the 33rd SIGCSE technical symposium on Computer science education* (pp. 23-27). New York: ACM. doi:10.1145/563517.563349

Mirkovic, J., & Benzel, T. (2012). Teaching cybersecurity with DeterLab. *IEEE Security & Privacy, 10*(1), 73-76. doi:10.1109/MSP.2012.23

Moore, J. (2011, August 15). Cyber recruits key part of NSA hiring blitz. *Federal News Radio*. Retrieved from http://www.federalnewsradio.com/?nid=241&sid=2497197

O'Connell, M. (2011, February 14). Budget request sees DHS increasing cybersecurity spending. *Federal News Radio*. Retrieved from http://www.federalnewsradio.com/?nid=473&sid=2747015

O'Leary, Mike. (2006). A laboratory based capstone course in computer security for undergraduates. In *Proceedings of the 37th SIGCSE technical symposium on Computer science education* (pp. 2-6). New York: ACM New York. doi:10.1145/1121341.1121346

Padman, V., & Memon, N. (2002, June). *Design of a virtual laboratory for information assurance education and research*. Paper presented at the 2002 IEEE Workshop on Information Assurance and United States Military Academy, West Point, NY.

Ragsdale, D. J., Lathrop, S. D., & Dodge, R. C. Jr. (2003). A virtual environment for IA education. In *IEEE Systems, Man and Cybernetics Society Information Assurance Workshop* (pp. 17-23). Piscataway, N.J: IEEE. doi:10.1109/SMCSIA.2003.1232395

Sadasivam, K., Samudrala, B., & Yang, T. A. (2005). Design of network security projects using honeypots. *Journal of Computing Sciences in College, 20*(4), 282-293. Retrieved from http://dl.acm.org/citation.cfm?id=1047890

Schafer, J., Ragsdale, D. J., Surdu, J. R., & Carver, C. A. (2001). The IWAR range: A laboratory for undergraduate information assurance education. *Journal of Computing Sciences in Colleges, 16*(4). Retrieved from http://dl.acm.org/citation.cfm?id=378720

Schweitzer, D., & Boleng, J. (2009). Designing web labs for teaching security concepts. *Journal of Computing Sciences in Colleges, 25*(2), 39-45. Retrieved from http://dl.acm.org/citation.cfm?id=1629042

Stackpole, B. (2008). The evolution of a virtualized laboratory environment. In *Proceedings of the 9th ACM SIGITE conference on Information technology education* (pp. 243-248). New York: ACM New York. doi:10.1145/1414558.1414618

Stewart, K. E., Humphries, J. W., & Andel, T. R. (2009). Developing a virtualization platform for courses in networking, systems administration, and cyber security education. In *Proceedings of the 2009 Spring Simulation Multiconference* (pp. 65:1-65:7). San Diego, CA: Society for Computer Simulation International. Retrieved from http://dl.acm.org/citation.cfm?id=1639877

Stockman, M. (2003). Creating remotely accessible "virtual networks" on a single PC to teach computer networking and operating systems. In *Proceedings of the 4th conference on Information technology curriculum* (pp. 67-71). New York: ACM New York. doi:10.1145/947121.947137

Tikekar, R. & Bacon, T. (2003). The challenges of designing lab exercises for a curriculum in computer security. *Journal of Computing Sciences in Colleges*, *18*(5), 175–183. Retrieved from http://dl.acm.org/citation.cfm?id=771860

Villanueva, B., & Cook, B. (2005). Providing students 24/7 virtual access and hands-on training using VMware GSX server. In *Proceedings of the 33rd annual ACM SIGUCCS fall conference* (pp. 421-425). New York: ACM New York. doi:10.1145/1099435.1099528

Wagner, P. J., & Wudi, J. M. (2004). Designing and implementing a cyberwar laboratory exercise for a computer security course. *Proceedings of the 35th SIGCSE technical symposium on Computer Science Education* (pp. 402-406). New York: ACM New York. doi:10.1145/1028174.971438

Wang, X., Hembroff, G. C., & Yedica, R. (2010). Using VMware vCenter lab manager in undergraduate education for system administration and network security. In *Proceedings of the 2010 ACM conference on Information technology education* (pp. 43-52). New York: ACM New York. doi:10.1145/1867651.1867665

Yang, T. A., Yue, K., Liaw, M., Collins, G., Venkatraman, J. T., Achar, S., & Chen, P. (2004). Design of a distributed computer security lab. *Journal of Computing Sciences in Colleges, 20*(1), 332-346. Retrieved from http://dl.acm.org/citation.cfm?id=1040274

Yurcik, W. & Doss, D. (2001). Different approaches in the teaching of information systems security. *Proceedings of the Information Systems Education Conference 2001,* 32-22. Retrieved from http://proc.isecon.org/2001/04a/index.html

Biography

Dr. Jason Pittman serves as a full time faculty member in the Information Assurance programs at Capitol College. Dr. Pittman's research focuses on game based learning in cybersecurity education, cyber competition design, gender and minority issues in cybersecurity instructional design, and artificial life models for cyber education. Prior to joining Capitol College, he worked in enterprise for more than 10 years as both a technical security implementer as well as in management and governance roles.