7-2012

Developing Small Team-Based Cyber Security Exercises

Brandon Mauer
Rochester Institute of Technology

Bill Stackpole
Rochester Institute of Technology

Daryl Johnson
Rochester Institute of Technology

Recommended Citation

Mauer B., Stackpole W., and Johnson D. Developing Small Team-based Cyber Security Exercises. In SAM'12 - The 2012 International Conference on Security and Management (Las Vegas, NV, USA, July 2012).

Developing Small Team-based Cyber Security Exercises

Brandon Mauer, William Stackpole, Daryl Johnson
Networking, Security, and Systems Administration
Rochester Institute of Technology
Rochester, New York
{bem5304,wrsics,dgjics}@rit.edu

Abstract—**The growth of the security industry is sparking a significant interest in well-rounded security professionals. Regional and national competitions in the academic community have been developed to help identify qualified candidates to support this industry. A course has been built to allow students to improve their skills in this area. This paper describes the process used to administer events in the support of such a competitive environment, and the process by which appropriate infrastructures are developed.**

Keywords-cyber security education, security competition; information security

I. Introduction

In September 2011, the authors led a seminar course entitled "Cyber Defense Techniques." This course places students into small teams for attack-defend events on small enterprise networks, mimicking the style of the National Collegiate Cyber Defense Competition [2]. A major benefit of providing such an environment is the ability to expose students to the operations of malicious users that commit digital crime, improving their skills at defending systems and networks from attack. In this paper, we discuss the specific roles and responsibilities of our positions in the course, elaborate on pertinent details concerning the seminar and its operation, provide interpretation of the results of the activity performed by the students, and discuss the lessons taken away from this unique experience.

II. Abbreviations and Acronyms

A. Abbreviations and Acronyms

Below is a list of terms used throughout this document, along with their meanings:

- CCDC: Collegiate Cyber Defense Competition. One of the United States' premier collegiate cyber security competitions, held annually, culminating in a national competition featuring the winning team from each of the ten U.S. regions in San Antonio, Texas. This event is the basis for the course layout and structure.
- Blue team: the team of students chosen to secure and defend a small enterprise infrastructure and to complete business tasks [1,5].
- Grey team: the team of students responsible for the development of blue team infrastructure and the creation, delegation, and assessment of business tasks. Normally, these two roles would be split into two teams, white for business tasks and black for development of infrastructure [1,4].
- Red team: the team of students who will be acting as penetration testers for a given event. The team's responsibilities include reconnaissance, vulnerability identification, infiltration, data theft, and sabotage, as directed by the grey team [1].
- Inject: a business task for the blue team to complete. Injects are not "mandatory," although failing to perform an inject resulted in no points being awarded. Each inject was given an independent maximum point value and was scored by the grey team after a predetermined time.
- Service check: an operation performed by the event scoring engine to assess the functional and correct operation of a network service based on specific grey team criteria, run at a predetermined interval. A check was only successful if it fully met all criteria. A separate check was used for each network service.
- Service uptime: The amount of checks assessed by the scoring engine to be successful, typically expressed as a percentage.
- NSSA: The Rochester Institute of Technology department of Networking, Security and Systems Administration [3].
- Cyber Defense Techniques: The name of the seminar course described in this paper. It may also be referred to as "the course" or "the seminar."

III. Roles and Responsibilities

A. Roles and Responsibilities

A course instructor for Cyber Defense Techniques (Johnson, Stackpole) is given five primary responsibilities:

- To provide overall direction for the course
- To lecture students on technologies and techniques relevant to each team role in the events

- To provide counsel, insight, or assistance to student teams as requested
- To ensure fairness of competition by removing bias, while providing a gradual increase in event difficulty
- To assess student performance to provide an academic grade for the course

A teaching assistant for Cyber Defense Techniques (Mauer) is given three primary responsibilities:

- To provide counsel, insight, or assistance to student teams as requested
- To ensure fairness of competition by removing bias, while providing a gradual increase in event difficulty
- To operate the scoring engine during each event

B. *Execution of Duties*

The authors have acquired experience at the CCDC on both blue and red teams. This experience proved to be invaluable in providing counsel to students when needed. Understanding the format of CCDC, its goals, and the type of challenges presented during the competitions helped eliminate confusion amongst grey teams as to what types of challenges were appropriate to build into the infrastructure, and streamlined the communication of ideas between the grey team and us. Although any team could freely ask questions, this was most often used by the grey teams to overcome their lack of experience in doing work in this area. The seminar is an advanced course; as such, assistance was only given to students to achieve their desired goals by providing direction and opinion.

The students were placed into specific teams to try to achieve three teams of equal measure. To maintain a high-quality event free of bias, each event infrastructure was analyzed to ensure no steps were taken to gain an advantage from a particular team setup. After the completion of each event, the next grey team was expected to take into consideration the lessons learned from the previous event to provide a more challenging infrastructure for this new event. Each event infrastructure was subject to approval before being put into use to verify the appropriate difficulty level was met.

The grey team was also asked to provide all of the necessary information to properly score a network service so that the scoring engine could be properly configured. A configuration file was then written containing this information, which would be processed by the scoring engine during its operation. The scoring engine used was commissioned for the 2006 National CCDC, written in C and Perl. Each service would be checked by the engine every three minutes. A successful check was recorded if the service provided the expected response as designated by the grey team. Any other response was considered a failure, as the checks were designed by the original programmers to simulate how an end user of a system would attempt to use the service being checked.

IV. COURSE STRUCTURE AND LAYOUT

A. *General Course Layout*

The course was divided into two major components: lectures and events. The first four weeks of the course were lectures that provided students an understanding of the basic components of each aspect of the event. Such topics included a primer on the use of the nmap network scanner and the Metasploit Framework for penetration testing, as well as illustrating techniques for securing core network services such as the Domain Name System, Internet webservers, and File Transfer Protocol. Students were also briefed on the roles and responsibilities of the grey team, with which many students did not have previous experience. The students were then placed into teams of four, one for each of the three roles needed in each event. The student teams would take a different role for each of the three events to experience red, blue, and grey team operations. Starting in the fifth week, the remaining six weeks were used for the completion of three event rounds.

B. *Event Layout*

Each event lasted for two consecutive course meetings for a total of three to four hours of activity time, depending on the setup and teardown time needed to return the lab environment to its original state. The events were conducted on an isolated network comprised of eight VMware ESX hypervisor computers, hosting all of the virtual machines the students would be using to attack, defend, and monitor the event. This virtual infrastructure was only in use during course hours to enforce a supervised and controlled competition environment. A preparatory meeting was held before the scheduled start of a given event. At this meeting, the grey team introduced the specifics of each role to its respective team for that event, including the type of infrastructure to be defended, the priorities the blue team should consider for defending the given infrastructure and network services, and priority of targets for the red team. The grey team was also given permission to provide unclear, misleading, or false information if they chose to do so.

The meeting immediately following the completion of the event was reserved for debriefing from the grey team; typical components discussed included any observations they recorded, score analysis of blue team performance on network service uptime and inject completion, and any recommendations they wished to provide. Red team and blue team members were also invited to share their observations and opinions.

In between each event was a week-long period that the grey team would use to consult with us to build the infrastructure for the next event. Some suggestions to the grey team included what services were appropriate in the type of scenario they envisioned to develop, as well as techniques or ideas on how to introduce hidden vulnerabilities into systems the blue team would be defending. Modeling the course after the style of CCDC allows the grey team to provide an infrastructure that is "broken"; systems may not have been fully patched and up to date, services could have been left misconfigured, and hosts may have had backdoors or other malware already installed on them. As previously stated, each event was conducted entirely in a location that had no Internet access, significantly hampering the capability to patch systems, which provided an additional challenge for both red and blue teams.

V. Comparison to Similar events

In the United States, cyber security exercises have existed since 2001, with the establishment of the U.S. Military Cyber Defense Exercise (CDX) [4, 6, 7, 8]. The success of such an exercise has accelerated the growth of additional exercises and projects in this area. Such documented exercises include a continuous, live, internal cyber defense activity at the University of Texas at Austin [9], an international capture-the-flag event created by the University of California at Santa Barbara [10], and a semester length graduate course at Texas A&M [8]. Each of these events represents a slightly varied approach to education-oriented cyber security. As the success of these three events continued, a steering committee was established to develop a new cyber defense competition, containing members from the Center for Infrastructure Assurance and Security (CIAS) at the University of Texas at San Antonio, the University of Texas at Austin, and Texas A&M [6]. The resulting competition, the CCDC, is the primary inspiration for the course described in this paper.

The seminar provides many of the same components offered at the CCDC and the exercises referenced above, but no two events are identical. The primary objective shared between each exercise is to provide the students interactive learning opportunities in realistic scenarios. As the course described in this paper is of finite duration, the objectives are more objective and quantifiable over the prescribed period than the event developed by the University of Texas at Austin. Similarly, a capture-the-flag component is not present at the CCDC or in this course. Therefore, the most direct comparison can be made to the CCDC and the semester course found at Texas A&M.

A full-time project by the CIAS at the University of Texas at San Antonio, the CCDC performs scoring, judging, infrastructure development and Red Team assessment using neutral or third party entities and sponsors. The course described in this paper is smaller in scale than the CCDC, and temporal and financial constraints have limited the capability for third party sponsorship and participation in this exercise. However, by allowing the students to take on those additional roles, the students gain a well-rounded and more thorough appreciation for the specific operation of each team as well as an understanding of the considerable effort needed to bring such an exercise to fruition. The professional Red Team, White Team, and Black Team present at the CCDC create an environment that provides student Blue Teams considerable opportunity to demonstrate their cyber defense skill set, and the amount of feedback the professionals can provide is significant. As proponents of a comprehensive security education, the authors propose that structuring the course as described above allows the students to obtain an equivalent amount of knowledge over the duration of the course in a broader scope, by participating as a member of each team present in a given event. The authors recognize the potential for unfair treatment between student teams in an event, as the students will eventually design an event as the Grey Team and participate as Red Team members at some point during the course. As stated above, grey team infrastructure and inject inspection is required before an event begins, and the Grey Team as well as the course instructors actively monitor the competitors throughout the duration of each event round. The authors hope that continued experience in this fashion will allow for a more robust solution to maximize fairness and the opportunity for students to learn in such an environment.

VI. Observations on Team Performance

The interpretation of the results is centric to the interactions with each grey team for several reasons. It was necessary to interface with them directly to ensure the event was running as they intended, as they would be monitoring both teams for the entirety of the event. Any team-specific problems requiring attention would be channeled through the grey team to provide the competing teams more time and to streamline the event as a whole. To ensure flawless operation of the scoring engine, it was necessary to monitor its behavior at all times to correct any issues before they could impact scoring mechanisms.

A. Event One

The experiences observed with each grey team were more varied than anticipated. Many hours were spent with the first grey team discussing their plans to develop a fictitious online medical facility, complete with a website and generated patient records hosted in an SQL database. This event would showcase three major components: government compliance and standards enforcement, misconfigured services, and the OpenSolaris operating system. With a medical company based in the United States, patient records are subject to HIPAA regulations. As a result, the grey team later asked the blue team to provide a compliance report in this regard. The second most prevalent feature of this event was that of misconfigured services; the primary DNS server in this infrastructure allowed all zone transfers and supported dynamic updates from any machine. An attacker who is familiar with manipulating DNS could find a wealth of information about systems on the network and could easily change DNS information to suit their needs.

The first grey team worked tirelessly toward developing a complete, eight node infrastructure with well thought out business injects, centering on system auditing, service improvements, policy, and HIPAA regulation. This first event would set a high bar for the remaining two, as the blue team visibly struggled with the foreign operating system and had a low inject completion rate.

B. Event Two

The second grey team, recognizing the high quality displayed by the previous team, set out to complete an even more ambitious task--to develop a twelve node server farm with separate administration machines. This idea was especially unique for two reasons: the departure from a standard small enterprise scenario and the inclusion of systems already “pre-hardened”. Competitions of this style typically feature a variety of systems, such as client workstations (that would be used by an employee) and servers running corporate services and storing company data. In contrast, the server farm implemented by the grey team was similar to a hosting company. In this scenario, nearly all of the operating systems and data residing on blue team servers would be owned by a customer, a significant departure from the small enterprise convention. The blue team was also given a guarantee that some of the systems provided to them by the "customer" have already been "hardened," although to what extent these systems were hardened was not divulged. Therefore, the new grey team was urged to perform security hardening techniques on some of the systems of their choosing so as not to overwhelm the second blue team with an increased amount of systems to protect. This course, while designed to be challenging to the students, is still only an academic exercise and preparatory course, and with an approximate event time of three to four hours, there simply would not be enough time for either team to fully explore the infrastructure and realize the maximum benefit from this particular exercise.

The second event saw the return of misconfigured services and strongly emphasized the adherence to established service level agreements and contracts the company had entered into with the customer. The grey team explained to the blue team that one "site" had an uptime requirement of 75%, the second at 50%, and the third at 25%. While these levels were simplified, it was effective at forcing the blue team to prioritize the handling of issues as they arose.

The full infrastructure was not deployed as planned. The grey team was only able to deliver nine systems by the start of the event, plus two administration machines and a company-wide pfSense software firewall. The red team was able to capitalize on this and explored further into the network more quickly due to poor configuration of the firewall and one of the Active Directory domain controllers. Nonetheless, there was clear indication that the blue team was aware of the ramifications of violating their service level agreement with the customer (a large point deduction). As a result, only one SLA was not met at the end of the event (50% SLA).

C. *Event Three*

For the final event, the last grey team chose to implement a small online casino. This scenario was established early on, but they were unsure of how to complete their environment. Two suggestions were offered: a Pluggable Authentication Module (PAM) configuration that would let anyone log in to the Linux systems with any password, or a sendmail email server that would execute arbitrary commands sent to it in email messages. The grey team eventually chose to implement a misconfigured PAM to complement a poorly secured web server.

During the event, the last grey team was very observant and quick to respond to both the red and the blue teams. In particular, when they noticed the blue team was not working cohesively, the grey team initiated a mock fire drill to help the blue team regroup and renew their efforts. This proved to be beneficial; at the end of event three, the blue team had the highest inject completion rate of all three events.

VII. Event Decision Making Process

Even though the students were responsible for developing the infrastructure for each event and writing injects, there was one component of each event that was featured as a result of the authors' collaboration. As a result, each event had one or two components that not only tied the event together, but represented a realistic component of a small enterprise to which the students may have needed more exposure.

A. *Event One - OpenSolaris and HIPAA Regulation*

Many of the courses taught in the NSSA department at RIT are taught using Red Hat Linux systems. While Red Hat and its derivatives are enterprise-friendly and capable operating systems, they do not constitute such a large portion of enterprise operating systems that exposing students to other operating systems would seem unnecessary. Solaris, then, seemed the most appropriate choice, due to its orientation to enterprise use, proliferation in technical environments, and its stability when virtualized on the ESX cluster used to host each event. Solaris is a complex and sophisticated operating system that include features offered by few others, and is built on a tested UNIX core platform. Although there was little opportunity to showcase the more advanced features, it was our hope that the exposure to the open-source (and free) OpenSolaris would be an eye-opening experience. Feedback from RIT alumni from the NSSA programs indicate it is evident that Solaris is still used frequently in many types of enterprises that hire administrators, network engineers, or security professionals.

It was also important that students understand that they may be in possession of, or responsible for maintaining the confidentiality of information and documents. Medical records are one such type of information. As medical records are continuing to be made available digitally, regulations such as HIPAA will be more significant than ever. Companies who do business in this industry are bound by law to uphold these requirements concerning digital patient records, while ensuring doctors and other medical employees, as well as patients, can access the appropriate medical records in accordance with their rights. It behooves the persons responsible for the storage, safety, and security of this information to be vigilant in their duties to safeguard this information.

B. *Event Two - Provider/Customer Model and Service Level Agreements*

We strongly advocated for the development of a server farm/hosting company scenario for the second event. Server hosting companies and cloud computing vendors have changed the landscape in which customers and enterprises do business. This scenario would also provide practical demonstration of the issues that hosting companies and cloud vendors face when providing network services for customers. Combining the aspects of honoring contractual obligations and service level agreements together would help students who plan on working for a service provider to understand the need to develop clearly defined, enforceable, and effective security policies and service level agreements. These documents apply to internal business as well as business between provider and customer. Due to the evident prioritization of the blue team response, meeting two out of three SLA requirements is a step in the right direction.

C. Event Three - Dysfunctional Authentication and Compliance, Revisited

In the final event, the grey team deployed the infrastructure to the blue team with deliberately misconfigured PAM for all of the Linux systems. As an online casino, this company could potentially be responsible for upholding Payment Card Industry (PCI) compliance for credit card transactions; a database breach could potentially result in the theft of a large volume of credit card information. An attacker can do all manner of malicious things to a system remotely, but the danger is even greater if the authentication mechanisms that provide most of the access control to a system do not function properly. This behavior was particularly difficult to find, as the system would often accept password changes and other control modifications, but they were not enforced. Although many students found it humorous that they could manipulate the system freely, the realization that this system was very poorly secured as a result of a simple misconfiguration resounded clearly.

VIII. Lessons Learned

We have taken away several important lessons from these exercises. The students expressed that they have noticeably improved their skills, an opinion also shared by the authors. The students more firmly understand the value of team skills, and the observations presented in this paper help us to continue to refine our technique and expand our knowledge of operating this type of environment in which the students compete and learn.

A. Student Improvement

Based on our observations, students were not only more attuned to finding vulnerabilities in systems (both from an offensive and defensive perspective), but also more easily able to engage the thought processes of incident response and system auditing to improve system security. This experience as a seminar course provides additional learning in areas not typically covered by core curriculum, and is a superb addition to academic credentials and provides a broader foundation for continued study in this area.

B. Teamwork and Interpersonal Skills

The students have noticeably demonstrated a deeper appreciation of the importance of teamwork in an environment such as this. CCDC is a team competition, and many enterprises have teams of people or departments who work together frequently; consequently, an employee who can work successfully in a team setting can be very valuable. It was evident during the events that teams had initial internal friction over leadership and coordination; fortunately, much of it had been addressed by the end of the course, but teams who could not overcome that challenge admitted that they encountered continued difficulties.

This effect was most profoundly demonstrated by the grey teams. While the students are capable of engaging in both red and blue team exercises in other controlled environments, to our knowledge, the grey team is a unique experience offered in an environment such as this. Understanding the differences in approach from competing in an event versus designing, building, administering, and scoring an event offers a deeper insight into the true goals of the event. The effects of this insight go full-circle; an effective grey team can build an infrastructure that sufficiently challenges the blue team and gives the red team opportunities to hone their skills. Effective scoring of injects requires the capability to quantify that the blue team has achieved the understanding necessary to properly complete assigned tasks. This, in turn, helps discourage teams from completing objectives simply to obtain the points each objective is worth.

C. Knowledge for Future Endeavors

Our involvement in CCDC in the past was beneficial in helping the students get the maximum benefit from the course, but this knowledge is only part of the solution. Understanding the spirit of CCDC, its goals, organization, and structure ensures the course follows the same path laid out by those responsible for CCDC; however, the competition is not run by specific guidelines that mandate certain systems, devices, or components to be present (or not present). Understanding how systems work when functioning properly, the way individual software components work together to produce a functioning system, and how changes to the system affect its operation (both seen and unseen) are many of the remaining pieces of this elaborate puzzle. Engaging in the development of such an exercise is a large undertaking, which can be improved upon by steadfast practice and learning from the experiences of others. As a result, we are confident that we can use this experience as a foundation for future endeavors in this area and improve the quality of the course as it matures.

[1] ReferencesNational Collegiate Cyber Defense Competition, CIAS http://nationalccdc.org

[2] G. Vigna, "Teaching hands-on network security: testbeds and live exercises," Journal of Information Warfare, vol 3, no 2. 8-25 2003

[3] Networking, Security, and Systems Administration, Rochester Institute of Technology http://nssa.rit.edu

[4] Wayne J. Schepens, John R. James, "Architecture of a cyber defense competition," IEEE International Conference on Systems, Man and Cybernetics, vol 5. 2003, pp 4300-4305.

[5] Gregory B. White, Dwayne Williams, "Collegiate cyber defense competitions," ISSA Journal, October 2005.

[6] Art Conklin, "Cyber defense competitions and information security education: an active learning solution for a capstone course," in *Proceedings of the 39th Hawaii International Conference on System Sciences*. 2006. Honolulu, HI.

[7] Wayne J. Schepens, Daniel J. Ragsdale, John R. Surdu, "The cyber defense exercise: an evaluation of the effectiveness of infromation assurance education" The Journal of Information Security, 2002. **1**(2).

[8] Lance J. Hoffman, Daniel J. Ragsdale, "Exploring a national cyber security exercise for colleges and universities" Report No. CSPRI-2004-08, The George Washington University, Report No. ITOC-TR-04001, *United States Military Academy*. 2004.

[9] G. Chamalese and A. Pridgen, "The success of the UT IEEE communications society," *Proc. 8th Colloquium for Information Systems Security Education*, 10 June 2004, pp. 9-12

[10] G. Vigna, "Teaching network security through live exercises," *Proc. 3rd Ann. World Conf. Information Security Education* (WISE 3), Kluwer Academic Publishers, 2003, pp. 3-18