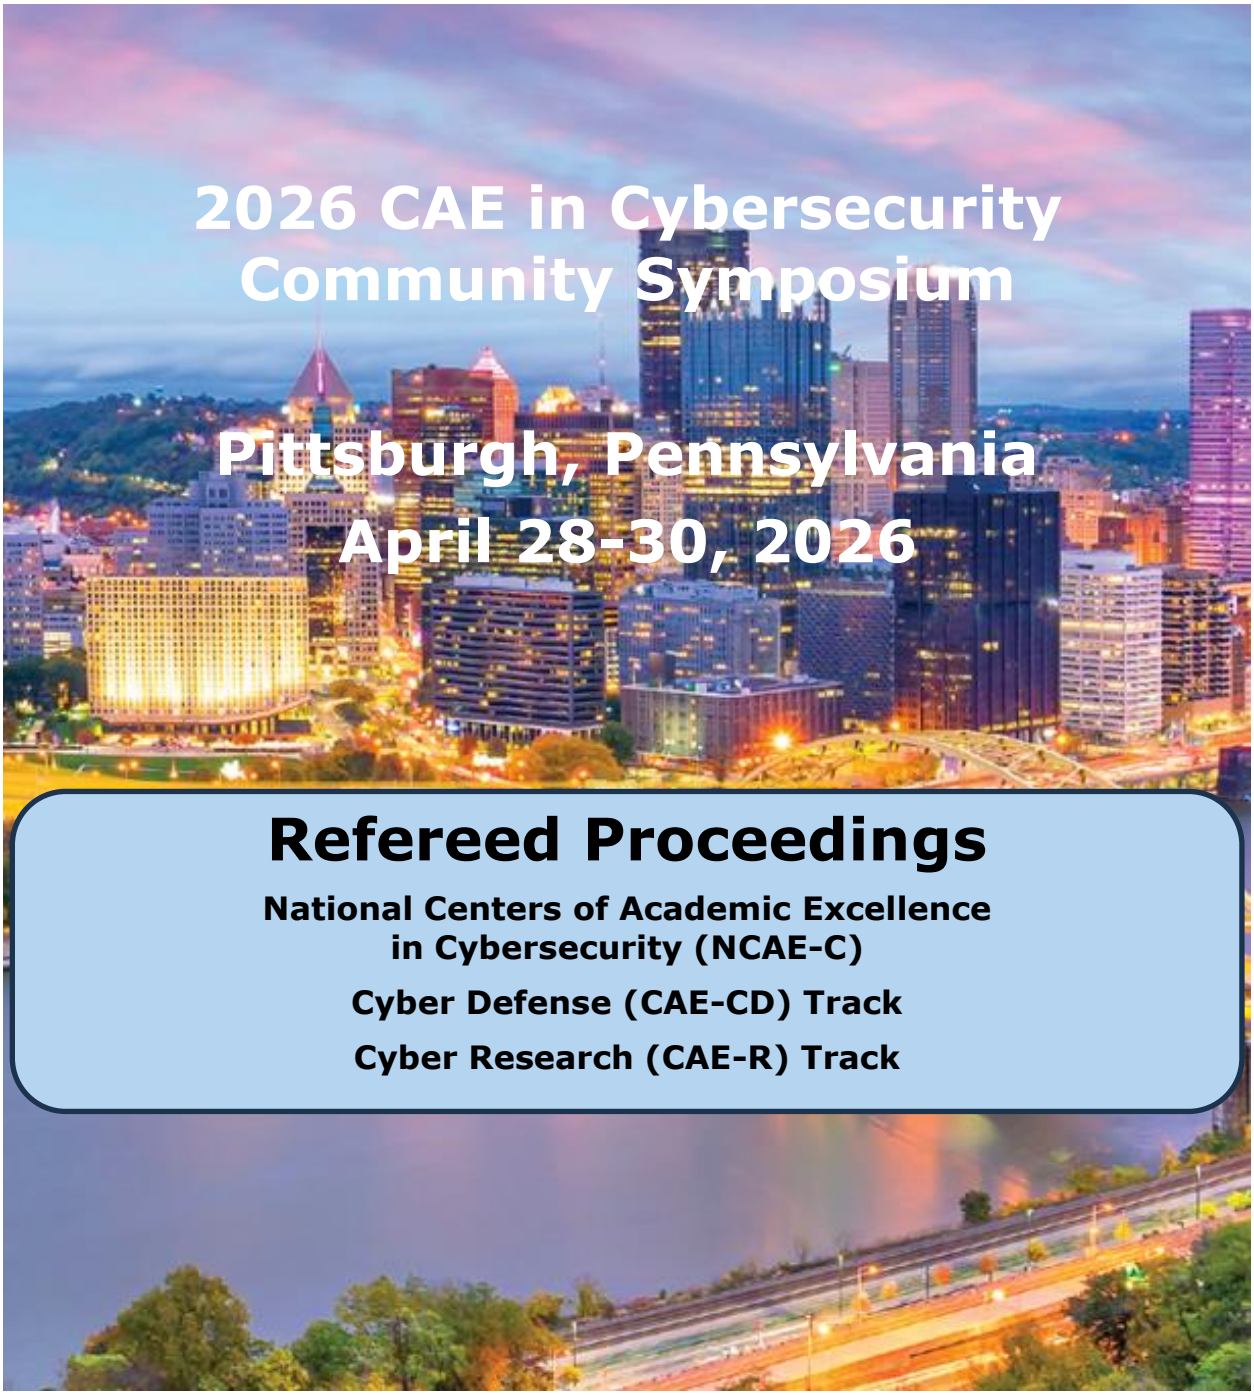

2026 CAE in Cybersecurity Community Symposium

Pittsburgh, Pennsylvania
April 28-30, 2026

Refereed Proceedings

National Centers of Academic Excellence in Cybersecurity (NCAE-C)

Cyber Defense (CAE-CD) Track

Cyber Research (CAE-R) Track

The CAE in Cybersecurity Community is funded by a National Centers of Academic Excellence in Cybersecurity grant HQ0034251E002 under the DoW Cyber Academic Engagement Office.

Table of Contents

CAE-CD CYBER DEFENSE

CD Track: Program Committee Co-Chairs

Tobi West
Coastline College, CA

twest20@coastline.edu

Yair Levy
Nova Southeastern University, FL
levyy@nova.edu

Anne Kohnke
University of Detroit Mercy, MI
kohnkean@udmercy.edu

CD Track: Program Committee Members

Name	Affiliation	State
Meryem Abouali	John Jay College of Criminal Justice	New York
Mohammed Abuhamad	Loyola University Chicago	Illinois
Goutham Reddy Alavalapati	University of Illinois	Illinois
Enas Albataineh	Florida Memorial University	Florida
Andrea Barrios	Coastline College	California
Debasis Bhattacharya	University of Hawaii Maui College	Hawaii
Gretchen Bliss	University of Colorado Colorado Springs	Colorado
Eric Burnett	California Institute of Applied Technology	California
Michael Burt	NCyTE	Washington
Prasad Calyam	University of Missouri	Missouri
Yuksel Celik	University at Albany SUNY	New York
Rohit Chadha	University of Missouri	Missouri
Eric Chan-Tin	Loyola University Chicago	Illinois
Kellep Charles	Capitol Technology University	Maryland
Ankur Chattopadhyay	Northern Kentucky University	Kentucky
Kristine Christensen	Moraine Valley Community College / NCyTE	Illinois
Art Conklin	University of Houston	Texas

Deanne Cranford-Wesley	North Carolina Central University	North Carolina
Ram Dantu	University of North Texas	Texas
Joan E. Debello	St. John's University	New York
Stuart Denrich	Stevenson university	Maryland
Thomas Devine	West Virginia University	West Virginia
Erdogan Dogdu	Angelo State University	Texas
Eman El-Sheikh	University of West Florida	Florida
Mohamed Elwakil	United States Coast Guard Academy	Connecticut
Nathan Evans	University of Denver	Colorado
Waleed Farag	Indiana University of Pennsylvania	Pennsylvania
Scott Fisher	New Jersey City University	New Jersey
Paige Flores	California State University San Bernardino	California
Kevin Floyd	Middle Georgia State University	Georgia
John Geiman	Western Dakota Technical College	South Dakota
Nicklaus Giacobe	The Pennsylvania State University	Pennsylvania
Sanjay Goel	University at Albany, SUNY	New York
Max Gorbachevsky	Utica University	New York
Katerina Goseva-Popstojanova	West Virginia University	West Virginia
Shwetha Gowdanakatte	Colorado State University	Colorado
Rob Greenberg	Sam Houston State University	Texas
Deniz Gurkan	Kent State University	Ohio
Jesse Hairston	The University of Alabama in Huntsville	Alabama
Jason Hammon	Western Governors University	Utah
Mohamed Hefeida	West Virginia University	West Virginia
Timothy M. Henry	Rhode Island College	Rhode Island
Ehren Hill	Virginia Tech	Virginia
Laura Hill	College of Western Idaho	Idaho
Robert Honomichl	University of Arizona	Arizona
Ann-Marie Horcher	Northwood University	Michigan
Md Tamjid Hossain	Texas A&M University-San Antonio	Texas
Randall Joyce	Murray State University	Kentucky
Jenny Ju	City University of Seattle	Washington
Justin Jutting	Western Dakota Technical College	South Dakota
Andrew Kalafut	Grand Valley State University	Michigan
Bilge Karabacak	University of North Carolina Wilmington	North Carolina
Siddharth Kaza	Towson University	Maryland
Tahir Khan	Western Illinois University	Illinois
Yoohwan Kim	University of Nevada, Las Vegas	Nevada
Gokhan Kul	University of Massachusetts Dartmouth	Massachusetts
Marufu Lamidi	Century College	Minnesota
Kevin Lann-Teubner	Butler Community College	Kansas
Mark Lawrence	New Mexico State University	New Mexico
Young Lee	Texas A&M University-San Antonio	Texas

Sandra Leiterman	University Arkansas Little Rock	Arkansas
Elliott Lynn	American Public University System	West Virginia
Olumide Malomo	Virginia State University	Virginia
Sebena Masline	Florida State College at Jacksonville	Florida
Mary McDonald	Anne Arundel Community College	Maryland
Suzanne Mello-Stark	Rhode Island College	Rhode Island
Stanley Mierzwa	Kean University	New Jersey
Jake Mihevc	Mohawk Valley Community College	New York
Stephen Miller	CAE Peer Review National Center	Washington
Jason Mitchell	Lansing Community College	Michigan
Janmejay Mohanty	South Carolina State University	South Carolina
Mousumi Munmun	Metropolitan State University	Minnesota
Van Nguyen	Saint Leo University	Florida
Laxima Niure Kandel	Embry–Riddle Aeronautical University	Florida
Cosmas Ifeanyi Nwakanma	West Virginia University	West Virginia
Aspen Olmsted	Wentworth Institute of Technology	Massachusetts
Loyce Pailen	University of Maryland Global Campus	Maryland
Glenn Papp	Niagara University	New York
Michael Ramage	Murray State University	Kentucky
Deep Ramanayake	Xavier university	Ohio
Douglas Rausch	Bellevue University	Nebraska
Indrajit Ray	Colorado State University	Colorado
Jeffrey Rice	Olivet Nazarene University	Illinois
Kevin Rickard	Moorpark College	California
Anna Rodgers-Stine	University of Alabama in Huntsville	Alabama
Suzanna Schmeelk	St. John's University	New York
Frederick Scholl	Quinnipiac University	Connecticut
Shamik Sengupta	University of Nevada, Reno	Nevada
Anthony Serapiglia	Saint Vincent College	Pennsylvania
Christian Servin	El Paso Community College	Texas
Chris Simpson	National University	California
Kenyatte Simuel	Ivy Tech Community College	Indiana
Mirco Speretta	Fairfield University	Connecticut
Stuart Steiner	Eastern Washington University	Washington
Alan Stines	Middle Georgia State University	Georgia
Geoff Stoker	University of North Carolina Wilmington	North Carolina
Nikunja Swain	South Carolina State University	South Carolina
Hondo Tamez	Johnson County Community College	Kansas
Cara Tang	Portland Community College	Oregon
Kasia Taylor	Anne Arundel Community College	Maryland
Diego Tibaquira	Miami Dade College	Florida
Mathew Heath Van Horn	Embry-Riddle Aeronautical University-Prescott	Arizona
Luis Vicente	Polytechnic University of Puerto Rico	Puerto Rico

Paul Wagner	University of Arizona	Arizona
Ping Wang	Robert Morris University	Pennsylvania
Melanie Williamson	Bluegrass Community and Technical College	Kentucky
Fan Wu	Tuskegee University	Alabama
Lei Xu	Kent State University	Ohio
Zhongmei Yao	University of Dayton	Ohio
Morgan Zantua	City University of Seattle	WA
Jason Zeller	Fort Hays State University	Kansas
Aeron Zentner	Coastline College	California
Kelei Zhang	Fort Hays State University	Kansas
Junjie Zhang	Wright State University	Ohio
Dmitry Zhdanov	Illinois State University	Illinois

CD Track: Proceedings Editorial Preface

2026 CAE in Cybersecurity Community Symposium
National Centers of Academic Excellence in Cybersecurity (NCAE-C)
Pittsburgh, PA

In 1999, the National Security Agency (NSA) launched the Center of Academic Excellence in Information Assurance Education program. Over the subsequent decades, we have witnessed a strategic evolution of the program to meet the burgeoning demands of the national security landscape. Today, the National Centers of Academic Excellence in Cybersecurity (NCAE-C) program stands as a premier collaborative effort, uniting academic institutions across the Nation to establish and maintain rigorous criteria for cybersecurity curricula. Nearly 500 institutions now hold designations in Cyber Defense (CAE-CD), Cyber Research (CAE-R), and Cyber Operations (CAE-CO). Reflecting the rapid shift in the technological frontier, the program has further expanded to include the recently established Cyber Artificial Intelligence (CyberAI) Program of Study Validation, ensuring our educational frameworks remain at the cutting edge of innovation.

At the direction of the NSA's Program Management Office (PMO), dedicated working groups of faculty members from the CAE Community and NSA subject matter experts have been collaborating to develop, document, and annually refine the criteria for the CD, R, CO, and CyberAI programs (also known as the Requirements Documents). This collective effort supports the necessary ongoing updates of the NCAE-C Knowledge Units (KUs), the maintenance of the CLARK platform for sharing high-quality curricula, and the infrastructure required for the designation and redesignation of member institutions. Furthermore, the CAE Community serves as a vital bridge between government agencies—including the Department of War (DoW) and the FBI, industry leaders, and educational partners. Together, we drive critical projects of national importance, ranging from K-12 outreach and veteran-focused certificate programs to the integration of Generative AI into cybersecurity defense and operations, all with a shared commitment to producing the resilient workforce required to protect our national infrastructure.

This year marks another significant milestone for the NCAE-C Community. Building on the momentum of last year's inaugural CAE Community Symposium – CD Track Proceedings Book, we further formalized our scholarly contributions by continuing the academic double-blind peer-review process via the EasyChair platform. This rigorous system was used to manage submissions for presentations, mini-workshops, lightning talks, and posters, culminating in our second published proceedings book. We remain committed to the value of collective knowledge—sharing best practices, successful collaborations, and solutions to community-wide challenges. This achievement was a massive undertaking, and we extend a special thank you to Dr. Tobi West, whose dedication and meticulous organization were instrumental in bringing these proceedings to life.

This year the Program Committee (PC) received a total of 137 submissions and accepted 68 Presentations, 34 Lightning Talks, (for a total of 28 hours of content) and 28 Posters. The PC members and co-chairs completed a total of 322 peer-reviews. We would like to thank all 115 PC

members listed in pp. 1-4 for their outstanding scholarly reviews and dedicated quality feedback to the authors.

The submissions for this year's CAE-CD Track reflect a pivotal shift toward the dual nature of Artificial Intelligence (AI) in the cybersecurity ecosystem, aligning with the national mandate established in Section 1514 of the FY26 NDAA, which calls for the development of robust cybersecurity educational programs and curricula specifically for AI. A significant majority of the abstracts focus on "AI for Security" and the "Security of AI," exploring how Large Language Models (LLMs) and Generative AI (GenAI) can be leveraged to automate Security Operations Centers (SOCs), build adaptive curricula, and enhance threat modeling. Conversely, several submissions address the critical need for "AI Red Teaming" and defending against adversarial machine learning, highlighting a community-wide urgency to enhance the security of the very AI tools being integrated into national infrastructure. Beyond AI, a strong secondary theme emerges around experiential learning and workforce "readiness" models. There is a notable emphasis on moving "beyond the textbook" through the implementation of student-led cyber clinics, virtual cyber ranges, and "Living & Learning" ecosystems. These submissions demonstrate a sophisticated approach by CAE-CD institutions to bridge the gap between academic theory and industry practice, particularly through experiential learning, micro-internships, apprenticeship models, and the use of Capture the Flag (CTF) platforms to translate classroom performance into verifiable, workforce-ready skills.

Finally, we would also like to thank Dr. Tony Coulson, Amy Hysell, Vanessa Zaldivar, and their team at the *CAE in Cybersecurity Community National Center (CCNC)* (https://caecommunity.org/) for supporting the *CAE Community of Practice in Cyber Defense (CoP-CD)* (https://www.caecommunity.org/cop-cyber-defense) and the 2026 CAE in Cybersecurity Community Symposium. We would also like to thank the DoW CIO – CAEO and NCAE-C PMO for the support of the program and the CoP-CD. It is an honor to serve this exceptional community of scholars, government, and industry professionals as we collaborate to strengthen our nation and programs in order to address the critical need to educate and build up the next generation of cybersecurity professionals.

The 2026 CAE Community Symposium - CD Track Program Committee Co-Chairs,

Tobi West, Ph.D.	***Yair Levy, Ph.D.***	***Anne Kohnke, Ph.D.***
Coastline College, CA	Nova Southeastern University, FL	University of Detroit Mercy, MI

CD Track: Refereed Extended Abstract Proceedings for Presentations

Adversarial AI: Breaking and defending machine learning

[Presentation]

George Meghabghab, Roane State Community College, TN,
meghabghagv@roanestate.edu

Extended Abstract

As artificial intelligence systems become increasingly integrated into critical infrastructure—from autonomous vehicles to medical diagnoses to financial systems—understanding their vulnerabilities has never been more important. This presentation explores adversarial machine learning, a rapidly evolving field that examines how attackers can manipulate AI systems through carefully crafted inputs and how defenders can build more robust models. Adversarial machine learning has emerged as one of the most critical challenges in AI security.

As models become more powerful and widely deployed, the attack surface expands dramatically. Recent incidents demonstrate that adversarial attacks are no longer theoretical; they work on production systems, can be executed physically (printed patches, modified signs), and represent a genuine threat to safety-critical applications. Attendees will gain hands-on insight into adversarial attacks through live demonstrations, learn about real-world security breaches, and understand the current state of defense mechanisms. This talk bridges theory and practice, featuring interactive code demonstrations that show how imperceptible changes to images can completely fool state-of-the-art neural networks. This presentation covers attacks (FGSM, PGD, C&W, DeepFool) that have been published in peer-reviewed venues (ICLR, CVPR, NeurIPS) and defense mechanisms actively used by industry leaders, including Google, Facebook, and OpenAI.

When selecting an adversarial attack method, speed and specific utility are primary considerations. The Fast Gradient Sign Method (FGSM) represents the fastest option available, making it the ideal choice for quick testing scenarios and educational purposes where immediate results are necessary. Slightly slower but highly effective is Projected Gradient Descent (PGD), which serves as the "gold standard" for evaluating the robustness of machine learning models. For tasks requiring a balance between performance and precision, DeepFool offers a moderate speed profile. Unlike the rapid testing offered by FGSM, DeepFool specializes in understanding decision boundaries and calculating minimal perturbations. This makes it particularly useful for researchers analyzing the precise geometric weaknesses of a model rather than just its general robustness.

On the slower end of the spectrum lies the Carlini & Wagner (C&W) attack, which prioritizes stealth and potency over speed. While it is less time-efficient than FGSM or PGD, C&W is the preferred method for breaking defenses and executing stealthy attacks. This trade-off highlights that while some methods are built for speed and standard evaluation, others, like C&W, are essential for penetrating hardened defenses.

Keywords: Adversarial machine learning, AI security, neural network robustness, FGSM, PGD, Carlini & Wagner Attack, DeepFool.

Metro State University Cyber Clinic: Pedagogical approaches, measured outcomes, and key insights

[Presentation]

Faisal Kaleem, Metro State University, MN, faisal.kaleem@metrostate.edu

Mousumi Munmun, Metro State University, MN, mousumi.munmun@metrostate.edu

Extended Abstract

Cybersecurity clinics represent a high-impact pedagogical model that addresses two persistent challenges in cyber defense education: meeting the cybersecurity needs of under-resourced organizations and preparing workforce-ready graduates. The Metro State University (MSU) Cyber Clinic, established in Spring 2024 as an NSA-sponsored initiative in collaboration with Minnesota IT Services (MN.IT), exemplifies how experiential learning can be systematically integrated into cybersecurity curricula to achieve measurable community impact (Cybersecurity Clinics Consortium, 2025).

The clinic teams, led by graduate students and supervised by faculty and state cybersecurity professionals, conduct comprehensive risk assessments for K–12 schools, nonprofit organizations, municipalities, tribal entities, and small businesses. All engagements follow the CIS Critical Security Controls Implementation Group 1 (IG1) framework, providing a standardized, scalable, and defensible assessment methodology.

Since its inception, the MSU Cyber Clinic has trained over 60 students and provided complimentary risk assessments to more than 35 organizations throughout Minnesota. Demand from clients has surpassed initial projections, with over 500 assessment requests submitted via public-sector partners. Feedback from clients consistently emphasizes the value of actionable and clearly communicated recommendations that would otherwise be financially unattainable. Students report substantial improvements in technical proficiency, client communication, and career readiness, supporting the clinic's effectiveness as a workforce development model.

This presentation will outline the clinic's pedagogical framework, operational structure, outcomes, and key lessons, with particular emphasis on its relevance to the CAE-CD Community. Topics will include faculty and government collaboration, student readiness requirements, quality assurance mechanisms, and scalability. Attendees will gain practical insights for replicating cybersecurity clinic models to enhance cyber defense education and promote community cyber resilience Metro State University Cyber Clinic (2025).

Keywords: Experiential cyber defense education, community cyber resilience, cybersecurity clinics, competency-based learning, workforce development.

References:

Metro State University Cyber Clinic. (2025). *Metro State Cyber Clinic.* https://www.metrostate.edu/mncyber/clinic

Cybersecurity Clinics Consortium. (2025). *Cybersecurity Clinics Consortium.* https://cybersecurityclinics.org

Research and development priorities for the maritime transportation system cybersecurity

[Presentation]

Ulku Clark, University of North Carolina Wilmington, NC, clarku@uncw.edu

Bilge Karabacak, University of North Carolina Wilmington, NC, karabacakb@uncw.edu

Geoff Stoker, University of North Carolina Wilmington, NC, stokerg@uncw.edu

Kasey Miller, University of North Carolina Wilmington, NC, millerk@uncw.edu

Jeff Cummings, University of North Carolina Wilmington, NC, cummingsj@uncw.edu

Edwin Garces, University of North Carolina Wilmington, NC, garcese@uncw.edu

Hosam Alamleh, University of North Carolina Wilmington, NC, alamleha@uncw.edu

Laavanya Rachakonda, University of North Carolina Wilmington, NC, rachakondal@uncw.edu

Extended Abstract

The Maritime Transportation System (MTS), a core component of U.S. critical infrastructure, faces increasing cyber risk due to legacy OT, heterogeneous vendor ecosystems, limited observability, and fragmented governance across ships, ports, and authorities.

This paper presents the MTS Cybersecurity Technology Roadmap (TRM), a year-long, expert-driven effort to identify priority R&D needs. Developed through multi-stage expert consultation, the TRM analyzes technology drivers, operational constraints, and capability gaps in the maritime domain. Using the NIST Cybersecurity Framework as an organizing structure, it identifies 44 R&D program areas and 153 key research questions addressing critical gaps in maritime CPS cybersecurity.

Several R&D themes emerge. First, gaps in IT/OT asset visibility and integrity validation hinder trusted baselines, motivating research in adaptive validation, automated asset management, and CPS mapping. Second, limited maritime-specific threat intelligence and telemetry constrain detection and response, highlighting needs in telemetry exchange, maritime honeynets, and shipboard IDS. Third, tight cyber-physical coupling creates poorly understood cascading risks, driving priorities in digital twins, MBSE, and quantitative risk assessment. Finally, governance challenges- such as inconsistent standards and fragmented authority -underscore the need for compliance automation tailored to maritime operations. The TRM provides a structured research agenda to guide coordinated R&D across academia, industry, and government, supporting scalable maritime cybersecurity innovation beyond ad hoc approaches.

Keywords: Maritime cybersecurity, OT security, cyber-physical systems, technology roadmap.

References:

Clark, U. (2025, September). *Maritime Transportation System (MTS) Cybersecurity Technology Roadmap (TRM)*. https://hdl.handle.net/20.500.14481/1433

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

AI-driven cyber resilience and security for drone video analytics operations in contested environments

[Presentation]

Prasad Calyam, Rohit Chadha, University of Missouri-Columbia, MO, calyamp@missouri.edu

Vijay Anand, University of Missouri-St. Louis, MO, vijay.anand@umsl.edu

Reshmi Mitra, Southeast Missouri State University, MO, rmitra@semo.edu

Extended Abstract

The Internet of Battlefield Things (IoBT) represents a rapidly evolving complex mesh of interconnected devices such as cameras, sensors, drones, rovers, etc. forming the backbone of modern military operations. In the IoBT realm, the Tactical Warfighting Edge (TWE) is defined as the forefront of military engagement environments where these IoBT edge devices operate as part of computation and communication tasks within a given mission. Users and devices need to share networks, systems, and data within the purview of specific missions on the battlefield. Ensuring the security and resilience of these drones that perform video sensing and analytics through ground control station (GCS) guidance, when they become hostile assets in adversarial hands remains a significant challenge. These concerns are further worsened by the constrained and contested nature of the TWE, where devices often operate under severe resource, power, and computational limitations. In such environments, even a small misstep in platform design or policy enforcement can result in catastrophic mission failures, where e.g., compromised drones or sensors may leak sensitive data, misidentify targets that can ultimately disrupt mission success.

In this presentation, we outline an AI-driven cybersecurity and resilience framework that synergizes software simulation and realistic hardware testbed components to validate drone guidance optimization in contested TWE network conditions. Specifically, a reinforcement learning based approach is used to guide single and multiple drones to continue a planned mission, in cases where network communication needs to be suspended to avoid detection by adversary or if the adversary uses tactics such as jamming, GPS spoofing or physical hijacking to disrupt the mission. By leveraging continuous monitoring, dynamic authentication, and fail-safe mitigation controls, the framework protects drone video analytics operations from attacks in TWE settings.

We detail a realistic testbed for both simulation and experimentation of use cases that were employed in Ukraine's Operation Spider's Web, where drone mission success had to be achieved under contested electromagnetic environments, as well as threat of hijacking, and network intermittence. We show mission resilience results of single drone and multi-drone simulations using a realistic grid environment with obstacles, no-fly zones, and dynamic threats. In addition, we show results of experiments with security components in physical testbeds with drones and industry-grade hardware to validate real-world performance. We conclude with details on on-going and planned collaborations between academia in the CAE-CD community and various DoD stakeholders for furthering both research and education initiatives related to counter UAS efforts.

Keywords: Drone mission security, tactical warfighting edge, AI-based resilience technique.

Why "Applied cyber defense" looks different at community colleges — and why that matters to CAE

[Presentation]

Kevin Rickard, Moorpark College, CA, krickard@vcccd.edu

Extended Abstract

Within the CAE in Cybersecurity – Cyber Defense (CAE-CD) community, applied cyber defense is often framed through four-year institutional models emphasizing upper-division coursework, extended research timelines, and capstone experiences. While effective, these models do not fully reflect how applied cyber defense is implemented at CAE-CD designated community colleges. Understanding these differences is important to advancing the CAE mission of developing a resilient and workforce-ready cybersecurity pipeline.

Community colleges serve diverse learners including first-generation students, career changers, and individuals seeking rapid entry into cybersecurity careers. These institutions operate under compressed academic timelines, strong employability expectations, and institutional IT constraints. As a result, applied cyber defense programs emphasize operational readiness, demonstrable competencies, and professional practices that can be validated within one- to two-year academic pathways.

This presentation examines how applied cyber defense is implemented at a CAE-CD designated community college through curriculum design, enterprise-style lab environments, cyber defense competitions, and structured operational experiences. An open-access cybersecurity laboratory with dozens of servers, routers, and switches allows students to build and troubleshoot systems while pursuing individualized projects. Competitions such as the National Cyber League (NCL) and the National Centers of Academic Excellence (NCAE) Cyber Games are incorporated into coursework to evaluate applied problem solving and operational readiness. On-campus internships place students in operational roles supporting network, endpoint, server, and cybersecurity administration.

Rather than replicating research-centric or capstone-heavy approaches common at four-year institutions, these programs emphasize realistic system operations, documentation, incident response workflows, risk awareness, and professional conduct as core cyber defense skills. Faculty roles extend beyond traditional instruction to include lab architecture, infrastructure management, and translation between institutional governance requirements and instructional practice.

Lessons learned include the effectiveness of competitions as authentic assessment mechanisms, the value of operational artifacts as evidence of competency, and the challenges of scaling applied environments while maintaining security and institutional compliance. Recognizing multiple models of applied cyber defense strengthens workforce alignment, expands participation in cybersecurity education, and supports scalable educational practices across institution types.

Keywords: Applied cyber defense, community college cybersecurity, cyber workforce development, cybersecurity education, CAE-CD.

Pioneering opportunities: Partnering computer science and autonomous equipment technician programs

[Presentation]

John Geiman, Western Dakota Technical College, SD, john.geiman@wdt.edu

Justin Jutting, Western Dakota Technical College, SD, justin.jutting@wdt.edu

Extended Abstract

As autonomous technologies reshape the mining and heavy equipment industries, post-secondary institutions must rethink how technicians are prepared for an increasingly convergent skill set. This presentation examines the development of a blended academic pathway that integrates the Autonomous Equipment Technician (AET) certificate with a two-year Computer Science degree, focusing on industry conversations that have identified opportunities through the merging of programs.

The AET certificate was developed in response to growing workforce needs and is believed to be among the first of its kind globally. Autonomous operations demand technicians who can work across mechanical, electrical, welding, and computer-based systems, skills that are not typically developed through single-discipline programs. Early attempts to source talent from adjacent fields proved insufficient, as industry partners emphasized the need for technicians able to maintain equipment while also supporting software, sensors, and data-driven systems.

To address this gap, faculty explored integrating the AET certificate with a two-year Computer Science degree to provide more in-depth skills. The resulting Technical Studies model allows students to earn an associate degree and a certificate, providing flexible pathways that align with industry expectations and support workforce advancement.

The blended programs have also expanded opportunities for manufacturer partnerships and applied learning. As autonomous systems generate increasing volumes of data, including the Computer Science curriculum assists in data evaluation, processing, and security, positioning graduates as both technical specialists and key contributors to system performance and safety. The presentation illustrates how integrated program design can support a future-focused approach to technical education while serving autonomous equipment industries.

Keywords: Autonomous equipment, self-driving.

Enabling realistic network attacks and defenses using a networked cyber range deployed on a single host

[Presentation]

Junjie Zhang, Wright State University, OH, junjie.zhang@wright.edu

Extended Abstract

Hands-on cybersecurity education requires students to engage with realistic networked systems where attacks and defenses produce authentic operational effects. However, many instructional cyber ranges either rely on simplified virtual labs that abstract away real network behavior or require multi-host and cloud-scale infrastructure that is costly and difficult to manage. This work presents the design and deployment of a real, fully networked cyber range implemented entirely on a single physical host, enabling high-fidelity cybersecurity training with minimal infrastructure overhead.

The proposed cyber range leverages GNS3 to integrate network emulation and virtualization to construct a multi-segment environment that preserves real Layer 2 and Layer 3 behavior, including switching, routing, and protocol interactions. Despite operating on a single host, the range supports interconnected attacker, victim, and infrastructure networks using real routers, operating systems, and services. Students interact with the environment using standard security tools and observe genuine packet flows, routing changes, and service impacts rather than simulated outcomes.

Within this environment, students conduct a wide range of realistic cyber operations, including ARP spoofing and man-in-the-middle attacks, routing and BGP-based attacks, denial-of-service attacks, and web application exploitation using intentionally vulnerable services. Because the cyber range maintains authentic network behavior, these attacks result in observable traffic redirection, service disruption, and protocol-level side effects that closely mirror real-world systems. The same environment also supports defensive activities such as traffic inspection, routing policy enforcement, firewall deployment, and service hardening.

A key contribution of this work is demonstrating that realistic, networked cyber training does not require distributed or cloud-based infrastructure. By leveraging open-source platforms, the cyber range is cost-effective, repeatable, and well suited for classroom deployment.

Keywords: Cyber range, cybersecurity education, network security, hands-on learning, network attacks.

Building an AI-cybersecurity ecosystem: CEROC's integrated approach to curriculum, experiential learning, and research

[Presentation]

Muhammad Ismail, Tennessee Tech University, TN, mismail@tntech.edu

Eric L. Brown, Tennessee Tech University, TN, elbrown@tntech.edu

Extended Abstract

Artificial intelligence (AI) is reshaping both cyber defense and cyber offense, increasing the need for cybersecurity programs to provide sustained, hands-on preparation for AI-enabled security practice. National emphasis on AI education and workforce development further reinforces this need (The White House, 2025; Wetzel, 2025).

The Cybersecurity Education, Research, and Outreach Center (CEROC) at Tennessee Tech University addresses this need through an integrated AI-cybersecurity ecosystem that combines curriculum, experiential learning, cyber range infrastructure, and research-informed instructional design. CEROC uses research and laboratory outcomes to continuously refresh instructional content and student experiences. In 2024–2025, CEROC supported over 150 students through coursework, informal learning, and hands-on training; 46 students participated in 10 cybersecurity competitions, earning first-place finishes in CPTC and CCDC regionals, second place in CyberForce, and first and fourth place in the InfoSec Nashville capture-the-flag competition. CEROC also launched its first AI-Assisted Cyber-Physical Security Competition, engaging 18 students across six teams in developing AI-based intrusion detection models using authentic data from smart power grid and drone swarm testbeds.

CEROC's infrastructure strengthens these educational interventions. In 2024–2025, the center expanded GPU clusters and advanced testbeds for quantum key distribution, smart manufacturing, cyber-physical power systems, and drone swarms. CEROC also activated $3.9 million in research funding across 23 successful proposals and produced more than 90 publications, showing how the research pipeline directly supplies new datasets, learning environments, and student-facing opportunities. This same ecosystem supported initiatives such as SHIELD, AI Corps, and Quantum Discovery Day, further connecting emerging technologies to workforce preparation. Additionally, CEROC launched its Student-Led Security Operations Center, which has employed over 15 students to date. The novelty of CEROC's approach lies in deliberately integrating curriculum, competition, cyber range infrastructure, and a research pipeline into a coherent educational strategy.

Keywords: Artificial intelligence, cybersecurity education, experiential learning, cyber range, cyber-physical systems, workforce development.

References:

The White House. (2025, April 23). *Fact sheet: President Donald J. Trump advances AI education for American youth*.

Wetzel, K. (2025, June 12). *The impact of artificial intelligence on the cybersecurity workforce*. National Institute of Standards and Technology.

Aligning governance, risk and compliance curriculum with government and industry workforce practitioner needs for employment

[Presentation]

Sandra Jetton Blanke, University of Dallas, TX, sblanke@udallas.edu

Matthew Davis, University of Dallas, TX, mbdavis_adj@udallas.edu

Extended Abstract

This presentation is designed to support professors teaching Governance, Risk, and Compliance (GRC) courses and to guide cybersecurity students aspiring to careers in GRC-related work roles. The presentation will be sharing information on student successes in being hired in GRC positions, student projects to increase competencies and the use of artificial intelligence (AI) within GRC positions.

Awareness of GRC has continued to grow since the "govern" function was elevated in the NIST Cybersecurity Framework (CSF) 2.0, from its prior placement within the *Identify* function to a foundational role that informs the five core cybersecurity functions: Identify, Protect, Detect, Respond, and Recover. This formally recognizes "govern" as an essential driver of effective cybersecurity across organizations. The CyberSeek workforce heat map indicates 356,000 openings within Oversight and Governance and 100,000 openings related to compliance and policy roles. This data highlights a significant and sustained demand for GRC-focused professionals.

Aligning GRC course content with government and industry needs is an important first step in preparing students for GRC roles. One role studied is the Department of Defense (DoD) Privacy Compliance Manager work role (DCWF Work Role EN-732) encompasses a broad complex set of knowledge, skills, abilities, and tasks. Given the practical constraints of academic course structures, the researchers engaged practitioners in PCM roles to assist in prioritizing tasks for integration into GRC related cybersecurity courses. Additionally, the practitioners emphasized the importance of soft or *power skills which can be developed while completing their degrees. The researchers find these skills are fre*quently underemphasized in technical curriculum.

By aligning GRC course content with practitioner-informed needs, priority KSAs, tasks and equally important power skills, cybersecurity students trained in GRC are more likely to become desired candidates within GRC positions.

Keywords: Governance, risk, compliance, Cybersecurity Framework (CSF) 2.0, CyberSeek, Department of Defense (DoD), Privacy Compliance Manager (DCWF Work Role EN-732).

From packet capture to copy-paste: AI in student submissions

[Presentation]

Stuart Denrich, Stevenson University, MD, sdenrich@stevenson.edu

Extended Abstract

Generative AI tools have become a productive "lab partner" to some students. This presentation takes a look at how students are using generative AI tools and image generators to submit work that looks technically impressive, somewhat accurate and refined—yet looks highly suspicious, usually out of scope of the assignment and most times incredibly wrong. This addresses the issue of how GenAI is shaping student work in light of academic integrity issues.

Using student examples, the session highlights some common giveaways: AI-generated screenshots, network diagrams that don't map to reality, packet captures that never existed, and written explanations way beyond the core concepts covered. When a simple follow-up question is asked the student fumbles. The goal isn't to shame students or play "gotcha," but to share experiences with other educators encountering similar situations.

The presentation includes examples of a practical "Multi-Question Method" designed to discourage AI-only submissions. Case study assignments are also highlighted where practical use of GenAI by students emphasizes critical thinking and validation through peer review. By requiring personal explanation, tool disclosure, documented struggles, annotated screenshots, and real lab evidence. This approach shifts students from copy-paste to engagement, and experiential learning without banning AI tools outright.

Takeaway: The time students spend using GenAI to do faux work could be better spent on developing the skillset to use GenAI as a Cybersecurity tool.

Keywords: Academic integrity, generative AI, cybersecurity education, assignment design, AI fakes.

From training to talent: Collaborative approaches to cybersecurity workforce development

[Presentation]

Eman El-Sheikh, University of West Florida, FL, eelsheikh@uwf.edu

Andrew Wright, University of Louisville, KY, andrew.wright@louisville.edu

Michael Tu, Purdue University Northwest, IN, michael.tu@pnw.edu

Adel Elmaghraby, University of Louisville, KY, adel@louisville.edu

Extended Abstract

Upskilling and reskilling programs are vital for preparing diverse learners to enter and advance within the cybersecurity field. This session highlights three innovative, NCAE-C–funded workforce development initiatives led by the University of West Florida (UWF), the University of Louisville (UofL), and Purdue University Northwest (PNW). Together, these programs demonstrate scalable, collaborative models for delivering flexible, skills-based cybersecurity and AI training. Presenters will share best practices, lessons learned, and reusable resources, concluding with a call to action for the NCAE-C community to replicate and adapt these successful approaches.

Led by UWF, the **CyberSkills2Work** program is a nationally recognized initiative launched in 2020. In partnership with seven NCAE-C institutions, it has delivered 85 flexible pathways aligned with over 30 NICE and DoD Cyber Workforce Framework (DCWF) roles. The program has trained more than 3,500 veterans, transitioning service members, and first responders, resulting in 4,500 digital badges and 17 industry certifications. Recognized by the White House, it is now expanding to include AI-enabled cybersecurity roles.

The **Pathways Coalition**, led by UofL, brought together 17 NCAE-C institutions to collaboratively develop and share cybersecurity curricula. The program provides online pathways for 10 DCWF roles, featuring online instruction, synchronous support and career coaching, and specialized modules in Cyber Risk Analysis, AI, Cloud, and Post-Quantum Cryptography. Learners earn digital credentials and receive vendor certification vouchers to accelerate their professional entry.

PNW leads the **Cybersecurity Workforce Certificate-based Training (CWCT)** program with six partner institutions. CWCT offers free online AI and cybersecurity certificates for military-affiliated learners and government employees. To date, the program has received over 17,000 applications and graduated more than 1,200 participants across various certificate levels. By showcasing these three distinct models, this session provides a comprehensive roadmap for building a robust, inclusive, and national-scale cybersecurity workforce.

Keywords: Cybersecurity, workforce, reskilling, AI, training, DCWF.

Developing the faculty pipeline: An international conversation on cybersecurity education

[Presentation]

Gretchen Bliss, University of Colorado Colorado Springs, CO, gbliss@uccs.edu

Michael Burt, NCyTE Center, WA, meburt53@gmail.com

Paige Flores, California State University, San Bernardino, CA, paige.zaleppa@csusb.edu

Extended Abstract

The global demand for cybersecurity professionals has created an urgent need to recruit and retain qualified cybersecurity educators across higher education institutions worldwide. While significant attention has been paid to workforce development, less focus has been placed on the critical challenge of building a pipeline of faculty who can prepare the next generation of cybersecurity professionals. This panel brings together perspectives from the United States, United Kingdom, Canada, and Australia to examine how different national contexts shape the experiences of new and early career faculty in cybersecurity education. By exploring these international perspectives, we aim to identify both common challenges and context-specific barriers that affect faculty recruitment and retention across diverse higher education systems.

The panel will examine existing support structures and initiatives designed to assist new and early career cybersecurity faculty in establishing successful careers. In the United States, the Centers of Academic Excellence in Cybersecurity (CAE-C) Cyber Defense Community of Practice has developed a New and Early Career Faculty Initiative that connects emerging educators with experienced mentors and community resources to strengthen the cybersecurity education pipeline. Similar challenges pertaining to faculty recruitment, professional development, and career progression are also prevalent in UK, Canada, and Australia. Panelists will share examples of programs, resources, and support mechanisms that have proven effective in their respective countries.

This panel will highlight strategies that have demonstrated success in supporting new and early career faculty across different international contexts. The session will conclude with recommendations for building stronger international networks and communities of practice that can provide sustained support for cybersecurity educators throughout their careers. By fostering dialogue across borders, we hope to strengthen the global cybersecurity education community and ensure that institutions worldwide can effectively recruit, develop, and retain the faculty needed to meet growing workforce demands.

Keywords: Faculty development, faculty recruitment, global perspectives, international collaboration.

Building the educators who build the workforce: The CAE-CD New and Early Career Faculty Initiative

[Presentation]

Gretchen Bliss, University of Colorado Colorado Springs, CO, gbliss@uccs.edu

Paige Flores, California State University, San Bernardino, CA, paige.zaleppa@csusb.edu

Extended Abstract

As the demand for skilled cybersecurity professionals continues to grow, building a strong pipeline of educators equipped to prepare the next generation is critical. The CAE-CD New and Early Career Faculty Initiative addresses this need by supporting emerging educators through virtual workshops, networking opportunities, and resource-sharing within the CAE in Cybersecurity Community. To date, the initiative has hosted 16 workshops, attracted over 340 registrants, and garnered more than 1,200 views on YouTube.

By focusing on new and early-career faculty development, this initiative aims to equip educators with the tools and knowledge needed to deliver high-quality cybersecurity education aligned with industry standards and workforce needs. This presentation provides an in-depth overview of the initiative's activities, highlighting successful strategies and key takeaways that have supported early-career faculty in their professional growth. Beyond reflecting on past efforts, the session will introduce new ideas for future workshops and professional development activities designed to address evolving challenges faced by cybersecurity educators. Proposed topics include innovative teaching methods, practical resources for the classroom, workforce framework alignment and opportunities to engage students through extracurricular activities.

Throughout the session, attendees will have multiple opportunities to provide feedback on presented strategies and contribute their own insights. Participants are invited to suggest additional topics or approaches that could further strengthen support for new and early-career faculty in the CAE in Cybersecurity Community.

Keywords: Faculty development, early-career faculty, professional development, workshops.

References:

National Centers of Academic Excellence in Cybersecurity. (n.d.). *CAE-CD new and early- career faculty members.* https://caecommunity.org/cae-cd-newand-early-career-faculty-members

YouTube. (n.d.). New and early-career faculty members playlist. YouTube. https://www.youtube.com/playlist?list=PLo3yqKgTfZlNHAYUJZUMDf8EN4R7FEzvl

Artificial intelligence as a core cybersecurity competency: Redefining skills for the modern cyber professional

[Presentation]

Elliott S. Lynn, American Public University System, WV, Elliott.lynn@mycampus.apus.edu

Extended Abstract

Artificial intelligence tools are becoming embedded across defensive operations, threat analysis, governance, and decision-making processes, reshaping expectations for the cybersecurity workforce. While cybersecurity education has traditionally emphasized technical depth in areas such as networking, systems security, and risk management, the ability to effectively leverage AI enabled tools is increasingly emerging as a foundational professional skill rather than an optional enhancement. Workforce surveys and employer feedback consistently highlight the need for graduates who can interpret AI generated outputs, validate evidence, and apply domain judgment in operational contexts.

This presentation examines artificial intelligence fluency as a core cybersecurity competency and explores how AI tool utilization intersects with existing technical, analytical, and ethical skill requirements for cyber professionals. Rather than positioning AI as a standalone topic or advanced specialization, the session frames AI literacy as an integrated capability that supports core cyber defense functions including threat detection, incident response, vulnerability assessment, and governance, risk, and compliance activities. To illustrate this integration, the presentation includes brief examples of how AI enabled capabilities such as alert enrichment, automated triage, and policy summarization map to common cybersecurity roles and learning outcomes.

Drawing on current workforce trends and emerging curricular practices, the presentation outlines a skills-based framework that connects AI enabled tasks to measurable learning objectives and assessment strategies. This framework highlights the importance of teaching students not only how to use AI tools, but how to evaluate outputs critically, apply domain knowledge to guide tool usage, and operate within ethical, legal, and policy constraints. The discussion also addresses implications for curriculum design, assessment integrity, and faculty development within CAE CD programs.

By reframing AI tool usage as an essential cybersecurity skill, this session contributes to ongoing conversations within the CAE CD Community regarding workforce readiness, curriculum modernization, and interdisciplinary integration. Attendees will gain a structured perspective on incorporating AI fluency into cybersecurity education in a manner that strengthens, rather than replaces, foundational cyber defense competencies.

Keywords: Cybersecurity workforce skills, artificial intelligence literacy, cyber defense education, curriculum design, professional competencies.

Can AI models assist in predicting potential red flags and identifying emerging risks in an organization?

[Presentation]

Enrique Pesantez, Fairfield University, CT, enrique.pesantez@student.fairfield.edu

Charles Shelley, Fairfield University, CT, charles.shelley@fairfield.edu

Akshay Mathur, Fairfield University, CT, amathur@fairfield.edu

Subhrajit Majumder, Fairfield University, CT, smajumder@fairfield.edu

Mirco Speretta, Fairfield University, CT, msperetta@fairfield.edu

Extended Abstract

Higher education institutions face a growing and increasingly complex landscape of cybersecurity risks that threaten their information assets, operational continuity, and institutional reputation. This study explores the use of Artificial Intelligence (AI), with a particular focus on Generative AI capabilities, to identify, analyze, and predict digital risks and potential red flags that affect institutions of higher education. The research evaluates the effectiveness of Generative AI in processing and correlating publicly available information from organizational websites, public records, and social media platforms to uncover cybersecurity, financial, and operational risk indicators.

The findings generated by the AI model are assessed to determine how accurately and comprehensively Generative AI can identify cybersecurity risks compared to traditional human analysis, as well as to examine the inherent limitations of relying solely on public data for organizational risk prediction. The red flags that were identified are then mapped to an industry standard cybersecurity framework (i.e., NIST) and associated families of controls to ensure alignment with established best practices and to identify remediation recommendations that are actionable. All collected information and analytical outputs are consolidated into a structured report intended for security managers, providing practical insights to support informed decision-making, risk mitigation, and support the governance of cybersecurity within higher education institutions. The proposed approach is further evaluated through the professional judgment of four subject matter experts (SMEs), who provided feedback on a rubric used to assess the quality and effectiveness of the AI-generated cybersecurity risk assessment reports.

Keywords: Generative AI, higher education cybersecurity, digital risk prediction, NIST Cybersecurity Framework, automated risk assessment.

Evaluating the impact of cybersecurity certifications on workforce readiness

[Presentation]

Waleed Farag, Indiana University of Pennsylvania, PA, farag@iup.edu

Extended Abstract

The growing reliance on interconnected computing systems in virtually every sector of modern life has significantly increased the need for robust cybersecurity measures to prevent compromise, disruption, and data breaches. This challenge is further exacerbated by a persistent global and national shortage of qualified cybersecurity professionals, a workforce gap that continues to widen as demand for skilled practitioners outpaces supply. To address this critical issue, numerous initiatives have emphasized cybersecurity education across all K–20 levels and encouraged the adoption of industry-recognized certifications as a means to equip students with essential, up-to-date skills. These certifications are widely regarded as a pathway to professional readiness; however, there is limited empirical evidence evaluating their actual effectiveness in preparing individuals for real-world cybersecurity roles. This study aims to fill that gap by analyzing the impact of certifications and complementary interventions on workforce development.

To achieve this objective, we conducted a comprehensive, systematic study examining the effectiveness of obtaining common cybersecurity industry certifications in preparing future professionals. The research explored how certifications and related interventions contribute to alleviating the national shortage of qualified cybersecurity experts. Our methodology employed a quantitative approach using both direct and indirect assessment techniques. Specifically, we designed surveys to measure key performance indicators such as certification attainment rates, perceived preparedness for cybersecurity roles, and job placement outcomes. Participants were recruited from 16 higher education institutions across Pennsylvania, resulting in more than 100 complete responses. The findings provide valuable insights into the role of certifications in shaping a skilled cybersecurity workforce.

This proposal presents the design, implementation, and key findings of the research study described above. Preliminary results indicate that while 80% of surveyed participants expressed interest in obtaining a cybersecurity certification, only 16% had attempted a certification exam. Of those attempts, 73% resulted in successful certification. Furthermore, 58% of certified participants reported feeling more confident in their professional abilities as a result of earning a certification, and 67% stated they would recommend their certification to others. These findings highlight both the strong interest in certifications and the gap between intent and action among students. In the full presentation, we will share additional insights derived from the study and discuss their implications for cybersecurity education and workforce development. We will also present conclusions drawn from in-depth statistical analyses of the collected data, offering evidence-based recommendations for educators, policymakers, and industry leaders seeking to strengthen the cybersecurity talent pipeline.

Keywords: Cybersecurity workforce gap, industry certifications, cybersecurity education, professional skill development, workforce readiness.

Designing a Student Security Operations Center to strengthen cyber defense workforce readiness at a community college

[Presentation]

Tobi West, Coastline College, CA, twest20@coastline.edu

Andrea Barrios, Coastline College, CA, abarrios16@cccd.edu

Extended Abstract

Cyber defense programs may struggle to give students meaningful experience in operational security roles, especially with limited resources. Security Operations Centers (SOCs) play a key role in cyber defense, yet most students do not have access to real SOC experience before entering the workforce. This paper describes the design, implementation, and early outcomes of a Student Security Operations Center (SSOC) embedded within a community college Computer and Cyber Sciences program to help address this gap.

The SSOC was developed as a structured, credit-bearing work experience aligned with the NICE Workforce Framework for Cybersecurity, with a particular focus on Cyber Defense-related work roles such as Defensive Cybersecurity, Incident Response, and Vulnerability Analysis. Students participate in a simulated SOC environment that emphasizes core cyber defense activities including monitoring and alert triage, incident investigation, threat analysis, documentation, and escalation procedures. Training is delivered through a combination of guided labs, bi-weekly synchronous sessions, and supervised operational tasks designed to mirror Tier 1 SOC analyst responsibilities without reliance on production enterprise data.

This project includes integration of the SSOC into the curriculum, including prerequisite coursework, competency-based learning outcomes, assessment strategies, and alignment to CAE Cyber Defense Knowledge Units. Particular attention is given to remote participation and accessibility, enabling participation by students balancing employment, family obligations, or geographic constraints. The SSOC model leverages simulation tools and structured playbooks to provide consistent, repeatable learning experiences while maintaining flexibility for institutional context.

Preliminary observations indicate improved student confidence in cyber defense concepts, strengthened knowledge of incident response processes, and increased readiness for internships and entry-level SOC roles. The model also supports equitable access to experiential learning by reducing barriers traditionally associated with on-site SOC environments. Early lessons from implementation included managing faculty workload, onboarding students effectively, and finding the right balance between instruction and realistic SOC operations. By presenting a practical, community college-based SSOC model, this project contributes a replicable approach for CAE-designated institutions seeking to strengthen cyber defense education through applied, workforce-aligned experiences. The findings demonstrate how student SOCs can serve as an effective bridge between classroom instruction and professional cyber defense practice while supporting CAE program objectives.

Keywords: Cybersecurity, workforce development, student security operations center.

Beyond the bio-lock: A behavior-based authentication system

[Presentation]

Mohammed Islam Tariqul, University of Illinois at Springfield, IL, misla80@uis.edu

Goutham Reddy Alavalapati, University of Illinois at Springfield, IL, goutham.ace@gmail.com

Extended Abstract

Smartphone authentication mechanisms primarily rely on point-of-entry verification to establish user identity at login time, after which the device remains unlocked until manually secured or a timeout occurs (Sitová et al., 2015). This design creates a vulnerability through which an unattended or compromised device can be accessed without further authentication (Shen et al., 2017). To address this limitation, we propose 'beyond the bio-lock', a continuous authentication framework that persistently verifies user identity throughout device usage. The system monitors subtle, natural user motion patterns—referred to as a behavioral heartbeat—that are continuously generated during interaction with the device (Shen et al., 2017; Yang et al., 2014). Any interruption or significant deviation in this behavioral signature triggers an active logout, immediately locking the device and thereby minimizing the risk of unauthorized access.

Beyond the bio-lock employs a self-supervised temporal density model based on an LSTM autoencoder, trained exclusively on data collected from the authorized user. The framework continuously acquires multimodal sensor data from the accelerometer, gyroscope, and touch events, capturing how the device is held, moved, and interacted with over short temporal windows of approximately 500 ms (Shen et al., 2017; Sitová et al., 2015; Yang et al., 2014). These individualized motion and interaction patterns are inherently difficult to imitate and can be effectively distinguished using multisensory fusion. A lightweight online learning component maintains a dynamic trust score that quantifies the similarity between current behavior and the enrolled user profile. While normal usage preserves a high trust score, anomalous scenarios—such as device handoff or placement on a surface—cause the score to decline rapidly, triggering an active logout when a predefined threshold is crossed.

Keywords: Smartphones, security, biometrics, authentication, machine learning.

References:

Sitová, Z., et al. (2015). HMOG: New behavioral biometric features for continuous authentication of smartphone users. *IEEE Transactions on Information Forensics and Security*, *11*(5), 877-892. https://arxiv.org/abs/1501.01199

Yang, Q., et al. (2014, November). A multimodal data set for evaluating continuous authentication performance in smartphones. In *Proceedings of the 12th ACM Conference on Embedded Network Sensor Systems* (pp. 358-359). https://doi.org/10.1145/2668332.2668366

Shen, C., et al. (2017). Performance analysis of multi-motion sensor behavior for active smartphone authentication. *IEEE Transactions on Information Forensics and Security*, *13*(1), 48-62. https://ieeexplore.ieee.org/abstract/document/8006292

Passwordless authentication

[Presentation]

Enas Albataineh, Florida Memorial University, FL, Enas.Albataineh@fmu.edu

Extended Abstract

The continued rise of credential-based cyberattacks highlights persistent weaknesses in traditional password-based authentication systems (National Institute of Standards and Technology (NIST), 2023). Despite extensive investments in password policies, user awareness training, and technical safeguards, passwords remain highly susceptible to reuse, phishing, brute-force attacks, and social engineering (NIST, 2023). Large-scale data breaches and increasing regulatory requirements have increased interest in alternatives to traditional password systems. Passwordless and adaptive authentication reduce reliance on shared secrets while strengthening access control and improving usability (FIDO Alliance, 2023).

This extended abstract examines contemporary authentication paradigms with a focus on passwordless authentication mechanisms, including biometrics, hardware security keys, and public key–based authentication (FIDO Alliance, 2023; NIST, 2023). The purpose of this work is to analyze how these approaches address long-standing password vulnerabilities while introducing new technical, behavioral, and organizational considerations. This abstract discusses how authentication factors based on possession and inherence significantly reduce attack surfaces associated with knowledge-based credentials (FIDO Alliance, 2023). In addition, the work explores how risk-based and continuous authentication models leverage contextual indicators, such as device posture, user behavior, and environmental signals, to dynamically strengthen authentication decisions (NIST, 2023).

This work synthesizes recent literature and industry standards on passwordless authentication. It evaluates its security benefits, usability, and deployment challenges, including issues such as user acceptance, biometric privacy, and organizational readiness (NIST, 2023). It also explains how passwordless systems support zero trust architectures, helping organizations reduce phishing and credential-based attacks while supporting compliance with modern identity and access management frameworks (NIST, 2023). This work provides a current overview of passwordless authentication as a secure alternative to traditional passwords (FIDO Alliance, 2023; NIST, 2023). It also offers practical insights for educators and cybersecurity practitioners to modernize access control systems.

Keywords: Passwordless authentication, biometrics, zero trust, and access control.

References:

FIDO Alliance. (2023). *FIDO2: Moving the world beyond passwords*. https://fidoalliance.org

National Institute of Standards and Technology. (2023). *Digital identity guidelines* (NIST Special Publication 800-63-4). U.S. Department of Commerce. https://www.nist.gov

Embedding community-based projects into IT and cybersecurity curriculum

[Presentation]

Denise Kinsey-Bergstrom, Franklin University, OH, denise.bergstrom@franklin.edu

Extended Abstract

There is often a disconnect between what is taught in the classroom and what students will encounter in cybersecurity positions in the real world. One way to align academics with reality is to infuse real world projects into the curriculum. Some can be simple tasks to complete while others can be course-long or program-long projects that immerse the student in the content, evidence competency, and offer a service to local organizations, thus improving their cybersecurity posture. These partnerships are the basis for many capstone courses, but the same principles can be embedded into any IT or cybersecurity course.

This presentation will highlight the insights, tools, techniques, and processes developed from over 40 successful community partnerships completed. Repeatable steps for making connections, assessing projects, scaling projects, managing communications and deliverables, and ensuring student participation and professionalism will be shared. Learn how you can build these opportunities into your programs without starting from scratch.

Key insights into scoping and scaling projects are shared to help determine if a project should be an assignment, a course-long experience, scaled over multiple courses and/or multiple terms. Experiences contributing to projects supported by multiple classes and different curriculum programs and faculty will be shared.

Keywords: Work-based learning, community partnerships, embedded curriculum, hands-on projects, portfolio building, artifact building.

Lessons learned from implementing undergraduate/graduate certificates in AI/ML and cybersecurity

[Presentation]

Luis M. Vicente, Polytechnic University of Puerto Rico, PR, luis.vicente@pupr.edu

Alfredo Cruz, Polytechnic University of Puerto Rico, PR, alcruz@pupr.edu

Extended Abstract

As the demand for cybersecurity and artificial intelligence (AI) skills grows, engineering programs face pressure to respond quickly while maintaining academic rigor and workforce alignment. Stackable credentials—structured pathways in which smaller microcredentials build toward broader certificate programs—have emerged as a practical mechanism to support professional growth and targeted technical specialization.

This extended abstract presents lessons learned from the design and implementation of undergraduate and graduate certificates in cybersecurity and AI/ML within an engineering department at a CAE-aligned institution. At the undergraduate level, microcredentials were embedded within existing courses and laboratories to provide focused competencies in AI applications and cybersecurity fundamentals, stacking into formal certificates. At the graduate level, advanced certificates deepen technical knowledge in AI/ML and cybersecurity through specialized coursework aligned with research-informed content and workforce demand.

Rather than emphasizing technical content, this work focuses on organizational and strategic dimensions of implementation. Key considerations included curriculum mapping to workforce-aligned competencies, cross-disciplinary faculty coordination, integration with existing degree programs, and grant-supported development to accelerate deployment. A central design principle was maintaining flexibility while ensuring defined learning outcomes and assessment standards consistent with institutional and CAE expectations.

Implementation challenges included balancing rapid development with academic approval processes, distributing faculty workload, and ensuring sustainability beyond initial funding. To address these challenges, strategies included modularizing existing courses, defining measurable learning outcomes, engaging faculty early in curriculum design, and aligning pathways with industry-informed competencies. These approaches supported feasibility and long-term viability.

The lessons presented aim to assist CAE community members developing credential-based pathways in cybersecurity and AI at both undergraduate and graduate levels, offering practical guidance for institutions seeking structured, workforce-aligned, and academically sound credential expansion.

Keywords: Cybersecurity education, artificial intelligence education, stackable credentials, microcredentials, workforce development, engineering education.

Bridging education and workforce: Hands-on cybersecurity with BCR and CWA

[Presentation]

Mary McDonald, Anne Arundel Community College, MD, memcdonald2@aacc.edu

Vini Nithianandam, Community College of Baltimore County, MD, vnithiana@ccbcmd.edu

Folashade Adeleke, Prince George's Community College, MD, adelekfo@pgcc.edu

Extended Abstract

In response to a rapidly expanding cybersecurity workforce gap, community colleges are leveraging advanced hardware and software platforms to provide hands-on training aligned with industry needs. A leading example of this approach is the Cyber Workforce Accelerator (CWA) program, a partnership between BCR Cyber and the Maryland Association of Community Colleges (MACC). CWA provides a sophisticated training platform that simulates cyber threats and defensive engagements in real time — to all 16 community colleges in Maryland.

This partnership represents a significant investment in specialized hardware and software, such as Splunk®, Palo Alto® and pfSense® that transcends traditional classroom instruction. By providing dedicated workstations connected to these cyber ranges, colleges now possess secure environments where students can interact with simulated networks, detect and respond to incidents, and apply theoretical concepts through practical exercises at no cost to them or the colleges. These capabilities address long-standing deficits in cybersecurity education, where limited access to expensive tools has constrained student readiness for workforce demands.

Hear from participating colleges about how this platform has catalyzed the revision and enhancement of cybersecurity curricula. Cyber range exercises have been integrated into digital forensics, ethical hacking and networking credit courses, ensuring that students graduate with applied technical competencies directly tied to employer expectations. Faculty are also involved in curriculum development, leveraging the range's simulations to reinforce competencies such as threat analysis, defensive architecture, and incident mitigation.

Future plans by MACC and BCR Cyber will foster cross-institutional collaboration, forming academic committees of faculty to share best practices and ensure that cyber range utilization continues to evolve with emerging threat landscapes and employer requirements. Thus, community colleges are not only recipients of technology but active partners in shaping the pedagogical frameworks that turn technology into workforce capacity.

The CWA framework also includes a workforce development pathway for BCR's SOC Operations Analyst I training and certification. In addition to sharing their best practices during this presentation, colleges participating in this initiative will discuss this pathway, its outcomes and challenges encountered and how they were addressed.

Keywords: Cyber range, cybersecurity, curriculum, design, instruction, hardware.

Cybersecurity curriculum at the speed of relevance: A growing national resource on CyberAI, quantum, and cyber operations

[Presentation]

Cara Tang, Portland Community College, OR, cara.tang@pcc.edu

Tobi West, Coastline College, CA, twest20@coastline.edu

Blair Taylor, Towson University, MD, btaylor@towson.edu

Sidd Kaza, Towson University, MD, skaza@towson.edu

Extended Abstract

Educators seeking high-quality, standards-aligned cybersecurity curriculum have access to a growing library of over 2000 free learning objects built by faculty on the CLARK library (www.clark.center). Developed through the NSA NCAE-C–funded Cybersecurity Curriculum Task Force (a partnership of 10 institutions led by Towson University, Portland Community College, and Coastline College), this national initiative offers open content in areas such as Reverse Engineering, Autonomous Vehicle Security, Zero Trust, Secure Software Development, Quantum Cryptography, Cyber Law, and more. Formed to create a scalable, sustainable model for cybersecurity curriculum development, the Cybersecurity Curriculum Task Force has completed curriculum reconnaissance, gap analysis, and production of high-impact instructional materials.

In 2024, Task Force 2.0 launched, with new content released or under development in emerging and critical areas including AI + Cybersecurity (Security of AI Systems, Applied Machine Learning, LLMs in Cyber), Secure Coding (Programming in Rust, OWASP Top Ten 2025), Quantum Computing (Quantum Principles, Risks in Encryption and Authentication, Quantum-Aware Cybersecurity), and Cyber Operations (AI and Cyber Operations, Programmable Logic).

Ensuring relevancy has been a foundational focus of the Curriculum Task Force. As the field of cybersecurity continues to evolve, driven by new adversarial tactics, technological shifts, and workforce needs, Task Force 2.0 is building new cutting-edge content and refreshing existing materials. All materials align to national frameworks, including the NSA CAE Knowledge Units (KUs), the DoD Cyber Workforce Framework (DCWF), and the NICE Workforce Framework, allowing institutions to integrate them directly into programs and support designation and accreditation efforts.

The work of this Task Force and the curriculum produced has been showcased at events including the 2025 CAE Symposium Curriculum Expo, CAE New & Early Faculty Meeting, CLARK demos, GenCyber events, CyberSkills to Work Coalition, SIGCITE 2025 Conference, and more.

Developed by faculty across the country, this growing collection of over 2,000 up-to-date learning objects represents one of the most comprehensive and free national resources for faculty seeking to expand or strengthen cybersecurity education at their institutions.

Keywords: Cybersecurity curriculum, CLARK, emerging topics, secure coding, artificial intelligence, quantum computing, cyber operations.

Securing AI-integrated healthcare data: Cybersecurity, ethics, and privacy

[Presentation]

Ourania Rahman, Fairleigh Dickinson University, NJ, o.rahman@student.fdu.edu

Harold Rosario, Fairleigh Dickinson University, NJ, h.rosario@student.fdu.edu

Camille Ellis, Fairleigh Dickinson University, NJ, c.ellis2@student.fdu.edu

Alexandre Kinkela, Fairleigh Dickinson University, NJ, a.kinkela@student.fdu.edu

Vansham Patel, Fairleigh Dickinson University, NJ, v.patel16@student.fdu.edu

Badi Aldousari, Fairleigh Dickinson University, NJ, b.aldousari@student.fdu.edu

Extended Abstract

Artificial intelligence (AI) is increasingly embedded within healthcare data infrastructures, transforming how electronic medical records, prescription data, and clinical decision-support systems are stored, processed, and analyzed. While AI improves operational efficiency, predictive accuracy, and clinical workflow coordination, it also significantly expands the healthcare cyberattack surface. The convergence of AI-driven analytics, cloud-based storage, wireless medical devices, and interconnected hospital networks exposes sensitive patient data to heightened risks, including ransomware, phishing, insider threats, and algorithmic misuse. This presentation examines the cybersecurity, ethical, and regulatory challenges associated with internally deployed healthcare AI systems and risks from external LLM tools through a security-first, cyber-defense lens. A structured framework is presented that maps healthcare data flows from collection and storage through analytics and automated decision support, identifying vulnerabilities at each stage of the data lifecycle. Key technical controls are discussed, including data classification by sensitivity, adaptive encryption based on real-time risk assessment, AI-assisted intrusion detection for anomaly monitoring, and secure authentication mechanisms for devices and users across hospital networks. Regulatory and ethical considerations are incorporated by comparing U.S. and European approaches to healthcare AI governance, focusing on the Health Insurance Portability and Accountability Act (HIPAA) and the General Data Protection Regulation (GDPR). The analysis highlights differences in transparency requirements, patient consent, automated decision-making oversight, and accountability structures, illustrating how regulatory gaps can introduce cybersecurity and ethical risk. Real-world incidents, including healthcare ransomware attacks and documented failures of clinical AI systems, are used to demonstrate the operational and patient-safety consequences of insufficient security and oversight. The presentation concludes with actionable recommendations emphasizing that AI can only safely advance healthcare when innovation is paired with robust cybersecurity architectures, ethical data governance, and continuous monitoring. By integrating technical safeguards with regulatory awareness, this work highlights practical strategies for protecting patient trust, healthcare system resilience, and public safety in AI-enabled clinical environments.

Keywords: AI in healthcare, cybersecurity, healthcare data protection, HIPAA, GDPR, ethical AI, cyber defense.

AI for security and security of AI in cyber education

[Presentation]

Dipankar Dasgupta, The University of Memphis, Memphis, TN, dasgupta@memphis.edu

Marco Carvalho, Florida Institute of Technology, Melbourne, FL, mcarvalho@fit.edu

Milos Manic, Virginia Commonwealth University, Richmond, VA, mmanic@vcu.edu

Extended Abstract

In the fast-evolving world of emerging technologies such as 5G/6G, Blockchain, IoT, edge Computing, autonomous vehicle/Drone, it is challenging to use intelligent decision-support systems (such as AI/ML techniques) since adversarial manipulations are possible and likely, as the attack surface has increased significantly with cyber-enabled interconnected systems and services. The multi-faceted AI/ML techniques appear to provide an efficient security paradigm to deal with influx of new threats in cyber-enabled critical infrastructures and various applications. These approaches can also be used to augment defense-in-depth and Zero-trust architectures and to provide necessary security enhancements to the design, implementation, and operation of legacy and future cyber-enabled systems.

Many cybersecurity solutions rely on black-box AI/ML tools that are poorly understood. At best, these applications contain blind spots that can be exploited, potentially leading to severe consequences. Unfortunately, many researchers are increasingly finding ways to insert backdoors into AI-based Large Language Models (LLMs) for defense, without a clear consideration or understanding of potential unintended effects. These malicious LLMs contain targeted functionality: Generate phishing emails, write polymorphic malware, automate reconnaissance.

The speakers will discuss the benefits and risks of using Generic LLMs in cybersecurity and LLM-based Agentic AI technologies. The workshop will also cover the topic on Security of AI/ML systems, since AI/ML based systems are also prone to adversarial attacks via input data manipulation and/or by modifying ML models. In addition, the organizers will address challenges of good vs. bad bias, industry vs academic approaches to AI in cybersecurity, moniker gaining recent popularity "there's no AI without security, and no cybersecurity without AI, as well as the role of data and what one should know to get "AI ready" data.

This workshop will highlight the importance of Cybersecurity research-based education and training in this new light of "AI for Security and Security of AI". The panelists will discuss the importance of introducing specific new topics, courses and approach to education curricula in this space. The CAE-CD institutes are engaged in cybersecurity education and outreach activities. The topics that will be covered in the workshop are very important in preparing next-generation cybersecurity workforce to deal with emerging technologies and associated security issues.

Keywords: Adversarial AI, AI security, cybersecurity education, critical infrastructure.

Securing AI systems through LLM red teaming: A practical pillar of modern AI governance

[Presentation]

Kellep A. Charles, Capitol Technology University, MD, kacharles@captechu.edu

Extended Abstract

As large language models and other advanced AI systems transition from experimentation into mission-critical environments, traditional security and compliance controls are no longer sufficient to address their unique risk profile. AI systems introduce new attack surfaces, including prompt injection, data leakage, model manipulation, and unpredictable outputs, requiring organizations to rethink how security and governance are operationalized.

This session presents a structured, practice-driven approach to securing AI systems through large language model red teaming as a core component of AI governance. The session is organized into three progressive modules: (1) foundational concepts and threat modeling for AI systems, (2) hands-on red teaming techniques, and (3) governance integration and operationalization. Participants will be introduced to established frameworks such as the NIST AI Risk Management Framework (MAP, MEASURE, MANAGE) and the OWASP Top 10 for LLMs to guide risk identification and testing strategies.

The practical component highlights specific techniques and tools used in LLM red teaming, including adversarial prompt design, jailbreak testing, and output evaluation using tools such as GARAK and DeepEval. Attendees will learn how to map testing activities to real-world risks such as sensitive information disclosure, model hallucination, bias, and system misuse. The session also demonstrates how red teaming outputs can be systematically documented and integrated into AI risk registers, impact assessments, and governance artifacts.

By positioning LLM red teaming as both a technical security function and a governance control, this session provides a repeatable methodology for embedding continuous testing into the AI lifecycle. Participants will leave with a clear understanding of how to structure red teaming efforts, select appropriate tools, and translate findings into actionable governance decisions that support safe, ethical, and resilient AI deployment.

Keywords: AI red teaming, AI governance, threat modeling, adversarial machine learning, risk management framework.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Digital risk mitigation for engineering majors

[Presentation]

Wm. Arthur Conklin, University of Houston, TX, waconklin@uh.edu

Extended Abstract

Engineers design and create the things that make modern society work, from highways to vehicles, from airplanes to buildings, factories, equipment, even all the things in data centers that enable Artificial Intelligence (AI) everywhere. Engineering is not a singular domain, but a set of related disciplines, each focused on different capabilities, such as electrical, mechanical, aerospace, industrial, etc. As the information age has changed society over the past 50 years, it has done the same to engineering. Today's engineering designs have significant information components with computers enhancing and enabling capabilities that we all rely upon. With the compute aspect comes networking and the need for security. Enter the world of digital risk mitigation. Engineers must determine the correct application of digital risk mitigation efforts to ensure that the system remains safe, reliable and efficient. The digital risk mitigation efforts, while similar to many in cybersecurity practice have many distinct differences, specifically a level of control that lies outside of the digital domain, making the mitigations stronger than typical cybersecurity efforts.

This presentation will introduce how digital risk has become embedded in engineered systems, and the need to control it via system design. A methodology developed by the U.S. Department of Energy through Idaho National Laboratory called Cyber-Informed Engineering, presents 12 principles that can guide engineers in developing systems that address the digital risk inherent in today's interconnected systems. The format is as follows:

1. Introduction to digital risk in engineered systems (Why SRE instead of CIA); how is security defined in an engineering system?
2. What is CIE in 3 minutes (elevator pitch version) How can this methodology help engineers understand digital risk and design mitigations into a system?
3. What are engineered controls and what is their role in mitigating digital risk in engineered systems
4. Why this topic belongs in CAE programs in engineering – why is this expansion of "cybersecurity" important and how is it different from IT based cybersecurity.
5. Where to go to get more information

The objective of this presentation is to provide CAE programs information on how digital risk can be addressed in an engineering program. Understanding the difference between IT cybersecurity and Engineered Systems digital risk mitigation efforts will facilitate the education and training of engineers and technicians with relevant skills to address the risks associated with interconnecting smart systems.

Keywords: Digital risk mitigation, engineered controls, Cyber-Informed Engineering (CIE).

From AI principles to campus practice: Lessons from implementing a knowledge-centered AI governance model in higher education

[Presentation]

Olumide Malomo, Virginia State University, VA, omalomo@vsu.edu

Ephrem Eyob, Virginia State University, VA, eeyob@vsu.edu

Extended Abstract

Artificial intelligence (AI) is increasingly embedded across higher education teaching, assessment, advising, and administrative operations. While national and international frameworks articulate ethical principles for responsible AI, many institutions struggle to translate these principles into consistent, campus-wide practice. As a result, AI adoption often outpaces governance, leading to fragmented policies, inconsistent academic-integrity expectations, privacy concerns, and limited institutional oversight. This challenge is particularly relevant for CAE-designated institutions, where cybersecurity, risk management, and compliance are central to institutional missions. This presentation examines AI governance through a knowledge-centered implementation lens, framing governance challenges not solely as technical or policy deficiencies but as organizational knowledge-management failures. Drawing on applied institutional experience and cross-institutional policy analysis, the session identifies recurring obstacles institutions encounter when operationalizing responsible AI use. These obstacles include ambiguous authorship and disclosure standards, inconsistent faculty guidance, limited AI literacy among students and staff, insufficient documentation of AI practices, and the absence of sustained monitoring and accountability mechanisms.

Rather than proposing new ethical principles, this session focuses on lessons learned from translating existing AI governance guidance into campus practice. It highlights how governance effectiveness improves when institutions treat AI oversight as an organizational learning process—one that requires deliberate documentation, shared understanding, structured communication, and continuous review. Emphasis is placed on the role of knowledge creation, retention, and transfer in ensuring that AI policies are not merely published but consistently understood and applied across academic and administrative units. The session further illustrates how weak knowledge-management practices contribute to governance breakdowns such as policy drift, uneven enforcement, and reliance on informal individual judgment. In contrast, institutions that align AI governance with knowledge-management principles are better positioned to support academic integrity, protect sensitive data, strengthen cybersecurity awareness, and maintain institutional trust. Designed for CAE-CD faculty, administrators, and program leaders, this presentation offers actionable insight into aligning AI use with institutional cybersecurity and compliance goals while navigating AI's dual role as both an innovation asset and a governance risk.

Keywords: Artificial intelligence governance, cyber defense education, "governance, risk, and compliance (GRC)", academic integrity, knowledge management, responsible AI.

Impact of GenAI in the classroom - Case study from a capstone course!

[Presentation]

Debasis Bhattacharya, University of Hawaii Maui College, HI, debasisb@hawaii.edu

Extended Abstract

Generative Artificial Intelligence (GenAI), which uses Large Language Models (LLMs), has existed for a few years and has significantly influenced higher education. Although educators understand what GenAI is, they find it challenging to implement it effectively in the classroom. While some disruptions have impacted student learning and assessment, they have also notably affected faculty teaching STEM, Business, and other disciplines.

This session will examine the impact on curriculum development, assessment strategies, and teaching methods for a business course, including a case study from a Business Capstone course in the ABIT BAS program at UH Maui College. Session learning outcomes will include designing an adaptive curriculum, project-based assessments, and oral student presentations that showcase creativity in solving new and unique problems.

Objectives of the presentation:

1. The presentation will provide a broad overview of GenAI and its latest advances
2. This presentation will examine the impact of GenAI on curriculum development, assessment strategies, and teaching methods
3. The presentation will include hands-on activities to engage a diverse audience

Keywords: Generative AI, capstone course, assessment strategies, pedagogy.

Adversarial thinking as an emerging professional disposition beyond computational thinking in cybersecurity education

[Presentation]

Christian Servin, El Paso Community College, TX, cservin1@epcc.edu

Nadia Karichev, El Paso Community College, TX, nmerzlya@epcc.edu

Ivan Alonso, El Paso Community College, TX, ialonso4@epcc.edu

Extended Abstract

Computational Thinking (CT) has long served as a foundational framework for introductory computing education, emphasizing decomposition, abstraction, pattern recognition, and algorithmic design. While essential, CT alone is insufficient for preparing students to design systems that are secure, resilient, and socially responsible. In cybersecurity education, learners must also develop Adversarial Thinking (AT), the capacity to anticipate misuse, failure, and exploitation, and to reason defensively about assumptions embedded in computational systems. This work positions adversarial thinking as an emerging professional disposition that extends beyond computational thinking and is critical for modern cybersecurity and computer science education. This work reports on a grant-sponsored initiative at a two-year program aimed at embedding adversarial thinking practices across introductory computing pathways. The project focuses on integrating cybersecurity-relevant reasoning into early computer science courses rather than isolating security into advanced or standalone classes. This approach aligns with national curricular guidance that frames security as a cross-cutting concern, including *ACM Computer Science Curricula 2023* and *Cybersecurity Curricula 2017,* and *NCAE Knowledge Areas/Units.*

The instructional model operationalizes adversarial thinking through repeatable pedagogical patterns: surfacing assumptions, modeling adversarial behavior, and designing mitigations such as validation, least privilege, and secure defaults. These patterns were infused into computer science fundamentals such as CS0, CS1, and CS2 using lightweight, faculty-friendly artifacts developed under the project, including adversarial misuse cases, secure coding checklists, and reflective assessments that reward defensive reasoning in addition to functional correctness. This work positions adversarial thinking as a necessary extension of computational thinking for preparing learners to design secure, resilient, and socially responsible systems.

Keywords: Adversarial thinking, secure and reliable programming, cyber fundamentals.

References:

Servin, C. (2025). Adversarial thinking as a professional disposition for computing education. *Proceedings of the 26th ACM Annual Conference on Cybersecurity & Information Technology Education (SIGCITE '25)* (pp. 161–167). Association for Computing Machinery. https://doi.org/10.1145/3769694.3771130

Capstones to clinics: Collaborative workforce

[Presentation]

Chris Simpson, National University, CA, csimpson@nu.edu

Teresa Macklin, California State University San Marcos, CA, macklin@csusm.edu

Extended Abstract

Three California universities have partnered with a regional cyber-focused nonprofit organization, to develop an innovative regional cybersecurity clinic model, supported by a $1 million grant from Google's Cybersecurity Clinics Fund. As one of the first multi-institution cybersecurity clinics in development, this collaborative effort demonstrates how academic partnerships can maximize impact and operational efficiency in serving underserved communities while providing students with hands-on cybersecurity experience.

The clinic model enables students from graduate and undergraduate programs in cybersecurity, homeland security, and computer science to work directly with small businesses, nonprofits, and community organizations to address real-world cybersecurity challenges. Under faculty supervision, students conduct security assessments, develop mitigation strategies, and provide practical recommendations, creating a win-win scenario where organizations receive needed cybersecurity support while students gain invaluable practical experience.

The presentation will share practical insights and lessons learned during the clinic's initial development phase, focusing on key challenges and solutions in establishing cross-institutional collaboration. Topics will include developing shared student training protocols, standardizing service delivery across institutions, coordinating resources, and managing centralized client intake processes. Presenters will discuss both successes and obstacles encountered while building partnerships between diverse academic programs and departments.

This session is particularly relevant to the CAE-CD community as it provides a practical roadmap for institutions considering similar collaborative ventures. The presenters will share their experiences in real-time, offering insights into the ongoing process of building a multi-institution clinic from the ground up. By highlighting both challenges and innovative solutions in establishing a regional cybersecurity clinic, this presentation will help other institutions develop more efficient and effective models for academic cybersecurity clinics serving their communities.

Keywords: Cybersecurity education, multi-institutional collaboration, community outreach, clinical model, service learning, regional partnership, capacity building.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Real-time proactive network intrusion detection via latent space optimization

[Presentation]

Kyle Wright, University at Albany, NY, kmwright@albany.edu

Lakshika Vaishnav, University at Albany, NY, lvaishnav@albany.edu

Sanjay Goel, University at Albany, NY, goel@albany.edu

Extended Abstract

Recent advances in Network Intrusion Detection Systems (NIDS) have emphasized early-stage detection, demonstrating that sequence-based models with attention mechanisms can identify malicious behavior before a network session completes. While effective, these architectures introduce substantial computational overhead. In high-throughput or resource-constrained environments, inference latency can approach or exceed session duration, limiting the practicality of proactive defense. This challenge is exacerbated by the growing volume and complexity of network traffic generated by cloud platforms, Internet of Things (IoT) devices, mobile systems, and industrial environments. Prior work has shown that much of this data is redundant for intrusion detection, motivating dimensionality reduction as a necessary step toward real-time operation.

This work proposes a framework that applies learned dimensionality reduction to packet-level network streams in order to reduce inference cost while preserving predictive utility. Autoencoders are used to compress high-dimensional representations into compact latent spaces through unsupervised training, enabling downstream models to operate on reduced inputs rather than raw packet data. The objective is not to replace sequence-based detection methods, but to make them computationally viable for early intervention by reducing the cost of processing evolving sessions. Using the IoT-23 dataset, packet-level streams are evaluated in both incomplete and complete session states to study early detection without reliance on finalized session information.

Dimensionality reduction is performed using multilayer perceptron (MLP)-based autoencoders, which project packet-derived representations into low-dimensional latent vectors. These latent vectors serve as inputs to attention-based temporal models trained to distinguish malicious from benign behavior as sessions progress. Evaluation focuses on standard classification metrics alongside system-oriented measures such as inference time and time-to-detection relative to session duration. The emphasis is on examining efficiency–effectiveness trade-offs rather than asserting performance gains. The anticipated contribution of this work is a practical pathway toward SOC-ready deep learning systems capable of operating within the temporal constraints of live network traffic. By combining packet-level early detection with autoencoder-based dimensionality reduction, this framework explores how proactive NIDS can be made more computationally tractable while preserving established modeling paradigms.

Keywords: Network intrusion detection, early detection, dimensionality reduction, autoencoders, IoT traffic.

Securing the car key fob: Fix the tech or hack the people?

[Presentation]

Ann-Marie Horcher, Northwood University, MI, horcheram@gmail.com

Extended Abstract

Sophisticated hackers have learned to intercept a key fob signal using commercially available and completely legal equipment (Magda & Payne, 2024). Though efforts to engineer a solution are in development, over 200 million cars on the road are at risk. Garfinkel's usable security principles state the necessity to provide "Good security now" with the technology currently available (Garfinkel, 2005). Designers must create the best security possible with existing technology.

This research approaches the car key fob risk from two research perspectives. One strategy to improve security is to hack human behavior. Humans are frequently the weak link in security (Lautaha, 2021). In this case, changing the behavior with a usable alternative could result in humans being the strong link. A class project was created by an advanced cybersecurity class to investigate this alternative.

Using car fob hacking equipment, a brief demonstration of the risk was created. The students used the demonstration to educate a snowball sample of thirty-seven visitors to the largest outdoor auto show in North America. After seeing the demonstration, and a possible solution to prevent the problem, the subjects were surveyed to determine what solutions they might adopt. The acceptability of the solutions was highly dependent on the usability, but the demo convinced over 80% of the subjects that some key fob protection was warranted.

The second focus was to examine technology solutions to the issues. Car manufacturers are aware of the risk and are developing new technologies to address the issue such as rolling codes (Parameswarath & Sikdar, 2022). Human behavior determines how soon solutions become viable. The risk is influenced by how often cars are replaced. Cars are a big-ticket expense and replaced less frequently than other technology. The average age of a vehicle is 12.6 years, and people typically keep a car for about eight years, making the tipping point for safety at least a decade away.

Keywords: Auto cybersecurity, key fob, rolling codes, usability.

References:

Garfinkel, S. L. (2005). *Design principles and patterns for computer systems that are simultaneously secure and usable.* Massachusetts Institute of Technology.

Lautaha, V. (2021). *Understanding individuals' IT security factors in smartphone usages: A threat avoidance point of view*. Capella University.

Magda, D., & Payne, B. R. (2024). RFID key fobs in vehicles: Unmasking vulnerabilities & strengthening security.

Parameswarath, R. P., & Sikdar, B. (2022). An authentication mechanism for remote keyless entry systems in cars to prevent replay and rolljam attacks. *Proceedings of the 2022 IEEE intelligent vehicles symposium (IV)*.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Hands-on high school cybersecurity education through university partnerships

[Presentation]

Andrew Kalafut, Grand Valley State University, MI, kalafuta@gvsu.edu

Samah Mansour, Grand Valley State University, MI, mansours@gvsu.edu

Extended Abstract

To increase student interest in cybersecurity and support development of the future cybersecurity workforce, it is beneficial to reach the students prior to their college years. The recent development of an Advanced Placement (AP) cybersecurity curriculum and exam demonstrates a growing interest nationally in making cybersecurity education available to students at a high school level. However, many high schools are likely to face challenges in meeting the demand for high school cybersecurity instruction, including technical and institutional policy constraints limiting students' access to meaningful hands-on cybersecurity learning experiences.

In response to these challenges, we have developed and piloted an outreach program in which a local high school cybersecurity class has visited our university monthly since the beginning of the 2025-2026 academic year. During these visits, students participate in hands-on cybersecurity activities aligned with the current AP Career Kickstart Cybersecurity pilot curriculum. Rather than relying on high school computer lab infrastructure, our model brings students to the university, where they can leverage our institutional infrastructure to engage with hands-on learning through tools such as Wireshark, Packet Tracer, and Nmap. This approach is motivated by the restrictive computing environments in place at many high schools, which often prohibit such tools. By bringing the students to our campus, we can provide a controlled lab environment where students can safely experiment.

The lessons delivered during each 2-hour visit are structured as a combination of short, focused lectures, demonstrations of the tools, and guided hands-on activities. The lectures reinforce the learning that took place in the high school classroom and ensure that the students have the necessary background information to understand what they are seeing in the hands-on tool. The demonstrations ensure students understand the tool's interface so they may proceed in the guided hands-on activities at their own pace. In addition to the technical skill development supported by these activities, this program also exposes high school students to the university environment. The high school students interact with university faculty and graduate students, exposing them to a diverse array of expertise they may not otherwise have access to. Such exposure may help attract students who may not have otherwise considered continuing to post-secondary education. In this presentation, we will describe the structure and implementation of this outreach program. We will discuss the challenges we have encountered along the way and provide practical advice for institutions interested in adopting a similar model.

Keywords: Outreach, K-12, education, awareness, networking.

Revisiting community Wi-Fi security through research replication: An undergraduate capstone pedagogy

[Presentation]

Jason Zeller, Fort Hays State University, KS, jlzeller@fhsu.edu

Caden Mayer, Fort Hays State University, KS, c_mayer2@mail.fhsu.edu

Kelei Zhang, Fort Hays State University, KS, k_zhang4@fhsu.edu

Extended Abstract

Hands-on, research-driven learning experiences are essential for preparing undergraduate students for careers in cybersecurity and network defense. This extended abstract describes a capstone-level instructional methodology in which an undergraduate student revisits and replicates a community Wi-Fi security assessment study originally conducted approximately sixteen years ago. The original project evaluated residential wireless security deployment practices using a vehicle-mounted data collection platform consisting of a laptop, GPS receiver, Wi-Fi interfaces, and routing equipment, with results geographically visualized using publicly available mapping tools (Bannister et al., 2009). That study demonstrated that undergraduates, when properly mentored, can conduct meaningful, ethical, and technically rigorous field research in Information Assurance. Motivated by personal interest in wireless networks and cybersecurity, the student independently reviewed the original publication and proposed a replication study using modern hardware and software, including passive multi-radio Wi-Fi monitoring, GPS mapping, and comparative data analysis following the original methodology. Rather than emphasizing results, this project is intentionally structured as a pedagogical exercise in research replication, comparative analysis, and applied cybersecurity practice. Students integrate theoretical knowledge from courses in wireless networking, network security, and cybersecurity ethics with practical skills such as tool selection, data collection design, passive reconnaissance, geospatial analysis, and responsible disclosure considerations. The replication framework encourages critical thinking by requiring students to assess how changes in technology, standards, and user behavior influence both methodology and interpretation. Additionally, students gain experience with research planning, documentation, and professional communication, reinforcing workforce-relevant competencies such as problem formulation, project management, and technical writing, with student reflections indicating improved confidence in applying cybersecurity concepts to real-world environments. The presentation will focus on course design, mentoring strategies, assessment methods, and lessons learned in guiding undergraduates through longitudinal research replication in wireless security.

Keywords: Wi-Fi networking education, experiential learning, student mentorship, capstone design.

References:

Bannister, M., Zeller, J., & Jiang, K. (2009). Colloquium for information systems security education. *Proceedings of the International Joint Conferences on Computer, information, and Systems Sciences, and engineering (CISSE), June 1-3, 2009* (pp. 95–101). Washington. https://cisse.info/e/archives/11-2009/12-papers/136-s04p03-2009.

The experiences of women and girls at the NCAE Cyber Games and cybersecurity identity

[Presentation]

Jake Mihevc, Mohawk Valley Community College, NY, jmihevc@mvcc.edu

Extended Abstract

The purpose of this presentation is to share the findings of a 2025 study that sought to understand the experiences of women and girls at the NCAE Cyber Games and how their experiences influenced the development of their sense of a cybersecurity identity. The study examined the experiences of four women and girls who participated in the NCAE Cyber Games through the lens of Carlone and Johnson's (2007) identity framework of competence, performance, and recognition. Review and analysis of the qualitative interview data identified common themes, including participants' self-definitions of success, tension within their experiences in cyber competition, the centrality of preparation, and the influence of mentors and colleagues.

The study found that the performance of cybersecurity tasks, the recognition of meaningful others, and the resulting feelings of competence were accessible to participants within their NCAE Cyber Games experiences. The study identified participant feelings and behaviors consistent with the presence of imposterism, stereotype threat, deficits in self-efficacy, and a defensive communication culture. The study also found that the participants chose to focus on learning and appreciated supportive communication that enhanced their learning. Participants that performed tasks, were recognized by meaningful others, and felt competent experienced an increase in feelings of cybersecurity, learner, and leadership identity consistent with Carlone and Johnson's identity framework. Refinements to the design and implementation of cybersecurity competitions in the areas of recruitment, structure of the preparation and competition experience, communication climate, behaviors and achievements that are celebrated, and narratives on competition may increase the efficacy of competitions for cybersecurity identity development among women and girls.

Keywords: Cybersecurity, competition, identity, women and girls.

References:

Carlone, H. B., & Johnson, A. (2007). Understanding the science experiences of successful women of color: Science identity as an analytic lens. *Journal of Research in Science Teaching, 44*(8), 1187-1218.

Responsible intelligence: Designing a career and ethics-focused artificial intelligence curriculum for high school students

[Presentation]

Jesse Hairston, University of Alabama in Huntsville, AL, jesse.hairston@uah.edu

Anna Rodgers-Stine, University of Alabama in Huntsville, AL, anna.rodgers@uah.edu

Extended Abstract

The Regions Investing in the Next Generation (RING) Cybersecurity Course, impacting over 35,000 students and 1,500 educators across the United States is being expanded to include an artificial intelligence (AI) course, RING AI, set to be released nationally in Summer 2026. RING AI is a full-year course mapped to work roles within the DoW Cyber Workforce Framework (DCWF) and the Centers of Academic Excellence (CAE) AI Foundational Knowledge Units. The full-year curriculum is segmented into multiple short courses – AI Fundamentals, AI Governance, Laws, & Ethics, and Machine Learning Fundamentals – that map to a distinct DCWF work role, helping teachers to integrate reasonable portions into their classes and encouraging students to explore in-demand career options. Throughout the course, students will engage with the AI concepts with a clear focus on careers within the defense sector and a foundation in the ethical development and use of AI. The RING AI curriculum is provided at no-cost to teachers and students across all 50 states, Puerto Rico, and Washington, D.C. through grants by the DoW, the NCAE-C program office located at the National Security Agency, and CAE institutions.

The presentation will provide an overview of the primary goals of each short course, highlighting the DCWF role associated with each. The presentation will include a sample curriculum lesson including instructional materials and hands-on labs that participants will be able to work through during or after the session. In addition to this, the presentation will discuss the curriculum development strategies utilized in developing the curriculum, with a specific focus on identifying complex computing topics and building materials to make these topics accessible to high school students without prior cybersecurity or AI instruction. The presentation will end with an opportunity for participants to register for access to the course when it is released in Summer 2026.

Keywords: Artificial intelligence, machine learning, K12 curriculum, STEM education.

Leveraging CISA and CIS resources to design cyber resilience exercises for critical infrastructure education and training

[Presentation]

Elizabeth Rasnick, University of West Florida, FL, erasnick@uwf.edu

Extended Abstract

Cyber threats against critical infrastructure sectors pose complex challenges that require not only technical defenses but also well-coordinated human response capabilities. To prepare future cybersecurity professionals and current practitioners alike, educators and program developers must move beyond theoretical knowledge to include practical, resilience-driven exercise design in curricula and training programs. This presentation outlines an approach for integrating publicly available resources from the Cybersecurity and Infrastructure Security Agency (CISA) and the Center for Internet Security (CIS) into structured cyber resilience exercises tailored for critical infrastructure contexts.

Participants will first be introduced to key frameworks and exercise templates published by CISA and CIS that are designed to simulate realistic threat scenarios and decision-making environments. Building on these foundations, we present how the well-established tabletop exercise format can be adapted to academic and workforce development settings, providing a flexible structure that enhances situational awareness, incident response competency, and cross-functional communication among stakeholders. A significant component of the session examines the role of AI-powered tools in augmenting exercise creation, scenario variation, real-time facilitation, and assessment, enabling more dynamic and scalable training experiences.

The session also emphasizes inclusive engagement strategies that foster collaboration between students, faculty, industry partners, and government practitioners. Through examples and actionable steps, attendees will gain insight into embedding these exercises into course modules, capstone projects, and community partnerships, ultimately strengthening cyber resilience education in CAE-designated institutions. Attendees will leave with both conceptual frameworks and practical resources to implement and assess customized cyber resilience exercises that align with national best practices and evolving infrastructure threats.

Keywords: Critical infrastructure, cyber resilience, tabletop exercises, incident response training, CISA resources, CIS resources, scenario-based learning.

References:

Cybersecurity and Infrastructure Security Agency. (n.d.). *CISA tabletop exercise packages*. U.S. Department of Homeland Security. https://www.cisa.gov/resources-tools/services/cisa-tabletop-exercise-packages

Center for Internet Security. (n.d.). *Tabletop Exercises (TTX)*. https://www.cisecurity.org/ms-isac/tabletop-exercises-ttx

Retention and engagement measurement in cyber labs with lower overhead and unique parameters

[Presentation]

Nicklaus A. Giacobe, Penn State University, PA, nxg13@psu.edu

James A. Mayberry, Penn State University, PA, mayberry@psu.edu

Dhrupad Joshi, Penn State University, PA, dhrupad@psu.edu

Rishikesh Galande, Penn State University, PA, rag5984@psu.edu

Extended Abstract

This work reports efforts in developing cyber lab exercises in a lightweight browser-only lab system. This presentation outlines three different efforts, designed to reduce overhead, implement individual lab parameters, and to measure student performance as well as human cognition for students who have completed the lab. Implementation of networking in browser-based virtual machines (BBVMs) is challenging because the operating system runs inside the web browser. In previous work (Giacobe & Ruff, 2025), BBVM use was limited to file system orientation. The current work adds TCP/IP capabilities to BBVMs on the Buildroot Linux distribution. This allows simple network tools for cyber attack and defense inside of the small footprint operating system that runs completely inside of a common web browser. Larger size distributions take too much time to launch and are therefore unsuitable. Other features of Buildroot are examined for future cyber labs using this platform. To make labs delivered through this mechanism meaningfully repeatable, individualized parameters yield variations of a lab experience that are unique per student. Pseudo-random parameters in components of a network security lab (in intrusion detection) include IP addresses, port numbers, packet data and other functions that create variations of experience which can be used as answers to assessment questions.

Finally, work-in-progress is reported in the assessment of student engagement and retention. Psuedo-randomly generated parameters allow for the creation of almost infinite combinations of unique versions. Auto-grading mechanisms were created to assess if the student correctly submits the expected answer. While this permits students multiple attempts and infinite repeats, the question remains, “How many is enough?” Both engagement and retention measurement mechanisms are proposed in this work. In summary, this work extends current evaluation of cybersecurity labs, where curricular assessment has been based on exposure alone, and evaluation is limited to self-reported subjective feedback. Measuring skill and ability performance appears to be unique.

Keywords: Browser-based virtual machines, cyber education, curricular assessment.

Building a statewide cybersecurity pipeline: A replicable ecosystem model for capacity, access, and sustainability

[Presentation]

Sandra Leiterman, University of Arkansas at Little Rock, AR, saleiterman@ualr.edu

Extended Abstract

CAEs in Cyber Defense institutions face persistent challenges related to faculty capacity, credit portability, and sustainable pathway development, all of which directly affect program scalability, regional access, and workforce alignment. As demand for K–16 cybersecurity education increases, a national shortage of qualified cybersecurity educators constrains CAE institutions' ability to expand offerings and maintain concurrent enrollment pathways (Anderson et al., 2024).

This presentation describes a statewide cybersecurity education ecosystem developed in (state) that addresses these challenges through shared credentials, cross-institutional collaboration, and alignment beginning in high school. The Cyber Learn Network (CLN) shifts from institution-centric program ownership to a coordinated network model in which a common set of cybersecurity Certificates of Proficiency and Technical Certificates are offered across participating institutions and stack directly into BS degrees at four universities. Unlike traditional 2+2 pathways, this structure supports multiple entry, exit, and reentry points (across institutions) while preserving academic rigor, credit portability, and workforce alignment.

The CLN was established through coordinated engagement with the University of Arkansas System and has since expanded to include private institutions. Statewide alignment was supported through collaboration with the Arkansas Department of Education computer science initiative and the Arkansas Department of Higher Education, ensuring that concurrent enrollment, credentialing, and transfer policies supported cyber defense pathways.

Early outcomes show concurrent cybersecurity courses active in high schools and initial cohorts from community colleges transferring into four-year cybersecurity programs with full credit articulation. Key obstacles addressed through this model include faculty shortages, inconsistent course equivalencies, transfer credit loss, and limited institutional reach. Addressing these barriers allows institutions to scale offerings without duplicating programs or overextending limited instructional resources. The CLN model demonstrates how shared governance, common credentials, and cross-institutional collaboration can be used by CAE institutions to expand cyber defense education capacity, mitigate faculty constraints, and build sustainable, statewide pipelines.

Keywords: K-16, cybersecurity ecosystems, faculty capacity, stackable credentials, cross-institutional enrollment.

References:

Anderson, K., Hammon, J., Morris, M., Vikayanon, K., Spragg, A., Giwamogorewa, O., Jirapanjavat, S., III, C., & Coombs, L. (2024). *Help wanted: Cybersecurity educators.* https://www.nist.gov/system/files/documents/2024/12/19/Cybersecurity%20Educators%20Wanted%20White%20Paper%20%28December%202024%29_508complant.pdf

Expanding and improving cyber defense education using private cloud and nested virtualization: A case study in network security

[Presentation]

Glenn Papp, Niagara University, NY, grp@niagara.edu

Petter Lovaas, Niagara University, NY, plovaas@niagara.edu

Extended Abstract

For smaller or less-resourced academic institutions, developing and building effective cybersecurity education without sufficient technical infrastructure can be challenging and can make CAE-CD validation feel out of reach. Hajny et al. (2021) classified 39 out of 60 cybersecurity competencies as technical in nature after assessing the gaps between work roles and required expertise in cybersecurity. For proper and safe administration, delivery of many CAE-CD knowledge units requires more complicated configurations than a typical lab machine (e.g., network forensics, penetration testing). To make this delivery possible, smaller or less-resourced academic institutions must often acquire some supporting technologies. Appropriate selection of these technologies is crucial as cybersecurity curriculum development and delivery capabilities depend heavily on these technologies. As a basic example, without robust virtualization capability, sufficient storage, and a permissible network, cybersecurity students are typically unable to use tools like network and vulnerability scanners as they would in industry, which inhibits students' ability to conduct experiments and research. Additionally, technical, bureaucratic, and financial decisions must be carefully considered, as errors—especially in technology acquisition, design, and configuration—could very well destroy the chances of building a successful cybersecurity program.

This research addresses these barriers by proposing a conceptual network design using either open-source or proprietary virtualization management platforms that smaller or less-resourced academic institutions can use to create their own private clouds and deliver virtual desktop infrastructure, including nested virtualization and many of the same configurable services as public cloud, to their students at a fraction of the cost of public cloud. The presentation will include a review of a graduate network security course's curriculum before and after the introduction of the private cloud infrastructure and discuss preliminary findings that suggest access to the infrastructure positively transformed outcomes, including analysis of student learning objectives, student certification post-course, and delivery and administrative costs. For future research, we propose testing whether adoption of private cloud for technology delivery in cybersecurity education programs at academic institutions correlates with higher student engagement, retention, and better outcomes versus public cloud adoption and no cloud usage.

Keywords: Cyber defense education, private cloud, nested virtualization.

References:

Hajny, J., Ricci, S., Piesarskas, E., Levillain, O., Galletta, L., & De Nicola, R. (2021). Framework, tools and good practices for cybersecurity curricula. *IEEE Access, 9*, 94723–94747. https://doi.org/10.1109/ACCESS.2021.3093952

Developing cybersecurity skills in health informatics through immersive VR learning

[Presentation]

Samah Mansour, Grand Valley State University, MI, mansours@gvsu.edu

Andrew Kalafut, Grand Valley State University, MI, kalafuta@gvsu

Sawesi Suhila, Grand Valley State University, MI, sawesis@gvsu.edu

Extended Abstract

As healthcare systems experience rapid digital transformation, the demand for professionals with applied expertise in health informatics and cybersecurity continues to increase. Despite this growing need, many graduate programs lack experiential learning opportunities that reflect real-world healthcare workflows and cybersecurity challenges. With the U.S. Bureau of Labor Statistics projecting a 17% growth in health informatics occupations, this study investigates whether a virtual reality (VR)–based instructional module can enhance students' problem-solving accuracy, task efficiency, and confidence in performing cybersecurity tasks relevant to health informatics practice.

A pilot study using a cluster-randomized crossover design was conducted with 60 graduate students enrolled in a Health Informatics program. Participants engaged in both VR simulation–based instruction and traditional instructional sessions covering key cybersecurity topics, including phishing detection, authentication mechanisms, and foundational cryptographic concepts. Pre- and post-intervention surveys were administered to assess changes in cybersecurity knowledge, task performance, and self-reported confidence. In addition, usability ratings and open-ended feedback were collected to evaluate student engagement, perceived learning value, and overall experiences with the VR environment.

The results revealed statistically significant improvements in both cybersecurity knowledge and self-reported confidence ($p < 0.05$) across four key domains following VR-based training. Although immersive scenarios involving decryption and threat identification were perceived as cognitively demanding, they produced substantial gains in learner confidence. Overall usability feedback was positive; however, approximately 32% of participants reported minor technical or spatial challenges while navigating the VR environment. Students consistently expressed a strong preference for the interactive and immersive nature of VR and indicated interest in future modules incorporating collaborative and healthcare-specific scenarios.

These findings suggest that VR-based instruction can significantly enhance applied cybersecurity knowledge, learner confidence, and engagement among health informatics students. The results further indicate that students initially exposed to traditional instruction benefited substantially from subsequent VR experiences. While VR is not intended to replace lecture-based teaching, it serves as a powerful complementary tool for developing practical decision-making skills and secure data workflow practices. Overall, this study provides evidence supporting VR as a scalable and experiential approach to cybersecurity education within digital health environments.

Keywords: Virtual reality, health, bioinformatics, simulation.

Robust network anomaly detection via self-attentive latent modeling

[Presentation]

Fangshi Zhou, University of Dayton, OH, zhouf4@udayton.edu

Tianming Zhao, University of Dayton, OH, tzhao1@udayton.edu

Grant Neeley, University of Dayton, OH, gneeley1@udayton.edu

Zhongmei Yao, University of Dayton, OH, zyao01@udayton.edu

Extended Abstract

We develop a framework for detecting network traffic anomalies that utilizes self-attention to extract complex latent features. Our model consists of *m* variational autoencoder (VAE) instances in parallel (e.g., *m*=5) (Zhou et al., 2024). Individual *m* VAE models capture latent layers, z_1 to z_m, that are connected via a self-attention module. When training the model, we feed only *normal* data to the model. Our model effectively captures latent features of normal data and can reconstruct the input with very *small* reconstruction errors (Mean Squared Errors). After training, we leverage reconstruction errors to detect anomalies, as anomalous inputs result in *high* reconstruction errors than the normal data. An *m* value between 8 and 10 often yields the best balance of detection performance and computational overhead. While using reconstruction errors to distinguish anomalies from normal data is classic, our parallel self-attentive latent model design is new.

We apply our model to detect anomalies in the KDD'99 dataset, the UNSW-NB15 dataset, the NSL-KDD dataset, the CSE-CIC-IDS2018 dataset, the CIC-DDoS2019 dataset, and the ToN-IoT dataset. Our model outperforms many existing methods in detecting anomalies in those datasets. For instance, the model performance on the UNSW-NB15 dataset is shown in Table 1.

Table 1. Performance comparison on UNSW-NB15.

Model	Accuracy	Precision	Recall	F1 score
NB + RF	76.04	76.00	83.40	76.80
AE + CNN	90.60	97.90	-	93.80
DNN	88.97	90.00	88.00	87.00
VAE + DNN	93.01	95.21	91.94	93.54
Our model	**99.45**	**99.64**	**99.55**	**99.60**

Keywords: Latent feature extraction, VAE, self-attention, anomaly detection, network traffic.

References:

Zhou, F., et al. (2024). A Parallel Gumbel-Softmax VAE Framework with Performance-Based Tuning. *Proceedings of the 27th European Conference on Artificial Intelligence (ECAI)*, 2024.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Developing a CAE Cyber AI–aligned program across multiple entry pathways at a Hispanic-Serving Institution

[Presentation]

Timothy M. Henry, Rhode Island College, thenry@ric.edu

Extended Abstract

As artificial intelligence continues to transform the cybersecurity landscape, Rhode Island College has embarked on an innovative initiative to develop a program meeting the requirements for the CAE Cyber AI designation. This presentation outlines our systematic approach to implementing the knowledge units for the "AI for Cybersecurity Program" (AICyber) while planning expansion to the "Security of AI Program" (SecureAI), and the unique challenges and opportunities we encountered in this process. The curriculum redesign was informed by extensive consultation with multiple advisory boards, the Rhode Island Governor's AI Task Force, and industry partners, ensuring our program aligns with both academic standards and workforce needs.

Our primary challenge was to ensure maximum student access to the program while maintaining rigorous CAE designation standards. We accept students from multiple pathways: (1) transfer students from the local community college with either associate degrees in computer science or cybersecurity, (2) students completing two years at Rhode Island College in either our cybersecurity, artificial intelligence, or computer science programs. Significantly, the local community college's cybersecurity associate degree holds CAE designation, creating both opportunities for alignment and challenges in ensuring seamless transfer credit articulation.

When we initiated the program design, the first two years of our cybersecurity and computer science programs were substantially different, with minimal curricular overlap. This fragmentation threatened to extend the time-to-graduation for transfer students or to force difficult choices among core competencies. Our approach required comprehensive curriculum restructuring to create common foundational pathways while preserving program-specific depth. We analyzed the CAE Cyber AI knowledge units and mapped them against existing courses in each of our programs, while carefully balancing the mathematics requirements for AI.

As a federally designated Hispanic Serving Institution (HSI) with a high proportion of first-generation college students, we bring unique perspectives to workforce development in emerging technology fields. Our student population often faces economic barriers and competing responsibilities, making efficient transfer pathways and clear degree progression essential for completing the program. The Cyber AI designation development process has strengthened our already close partnership with the community college, establishing a model for community college-to-baccalaureate pathways other institutions serving similar populations might replicate. Our presentation will conclude with lessons learned, ongoing challenges, and recommendations for other institutions navigating similar programmatic development in resource-constrained, transfer-heavy, minority-serving contexts.

Keywords: CyberAI, AI, community college, HSI, curriculum, mathematics, public university.

A CTF testbed for cyber-physical systems

[Presentation]

Alfa Nyandoro, Associate Professor, Regent University, VA, anyandoro@regent.edu

Trenten Podach, Research Associate, Regent University, VA, trenpod@regent.edu

Faustino Kuvaoga, Research Student, Regent University, VA, fauskuv@mail.regent.edu

Extended Abstract

Modern critical infrastructure operates at the convergence of Information Technology (IT) and Operational Technology (OT), yet cybersecurity education rarely addresses the unique challenges of this integration. This paper presents a hybrid cybersecurity testbed that integrates traditional IT infrastructure with Allen-Bradley Micro820 PLC-based cyber-physical systems for undergraduate education deployed and tested through a Capture the Flag (CTF) competition. Our dual-platform architecture requires blue teams to operate, monitor and defend industrial controllers while simultaneously providing cybersecurity services to conventional enterprise systems (Windows/Linux servers, databases, etc.). This approach addresses the reality that attacks on industrial systems can originate from compromised enterprise networks, exploiting trust relationships between business systems and control platforms. Likewise, due to their limited security capabilities, industrial control systems can be an effective attack avenue for bad actors targeting IT systems. The testbed architecture deliberately exposes realistic integration patterns, i.e. engineering workstations running Connected Components Workbench reside on an IT network, creating authentic attack paths from enterprise compromise to control system manipulation via cyber-physical networking protocols.

Prior to participating in the competition, participants are introduced to OT by training using OpenPLC. Subsequently, participants – grouped in blue teams, will be tasked with balancing IT security best practices with OT operational requirements—decisions with immediate physical consequences when controllers manage critical infrastructure such as motor systems, batch operations, and temperature regulation are disrupted. The Micro820 controllers exemplify real-world constraints: no encrypted communications, minimal authentication mechanisms and limited processing power prohibiting security agents or advanced logging. In addition, IT systems can be patched during maintenance windows, whereas PLCs control continuous processes where downtime may mean production losses or safety risks. Students experience this tension firsthand as they receive simulated "business requests" requiring system modifications during ongoing red team attacks, forcing prioritization decisions between security and operational availability.

Our work emphasizes the importance of cross-function expertise on dual responsibility platforms incorporating IT and OT technology, thus underlining the essence of teamwork and showcasing the inevitability of operational pressures during security incidents. This helps to reveal limitations of theoretical strategies under real-world constraints. The topology reflects industry practice, i.e. a shared infrastructure with multiple vendors, which can be replicated by other CAE institutions.

Keywords: Cyber-physical systems CTF, Industrial control systems, IT/OT convergence, Programmable logic controllers.

Lessons from a persistent student-led cyber clinic at UNLV

[Presentation]

Yoohwan Kim, University of Nevada Las Vegas, NV, yoohwan.kim@unlv.edu

Juyeon Jo, University of Nevada Las Vegas, NV, juyeon.jo@unlv.edu

Extended Abstract

The UNLV Cyber Clinic was founded with a dual mission: to safeguard and educate small businesses while establishing a rigorous professional foundation for the next generation of cybersecurity professionals. Today, it operates as a thriving, year-round, student-led organization with more than 120 active members, including high school students and alumni. Since its founding the Clinic has delivered cybersecurity education to hundreds of community members and conducted outreach and assessments for over 145 local small businesses. The Cyber Clinic bridges the gap between academic theory and real-world cybersecurity practice. Students develop advanced technical competencies through structured internal training workshops, industry certification preparation, and participation in Capture the Flag (CTF) competitions. Simultaneously, they cultivate essential professional skills through industry conferences, networking engagements, and community-facing presentations.

These efforts provide small businesses with high-quality cybersecurity assessments and guidance at no cost. Central to the Clinic's success is the immersion of members in real-world client acquisition, engagement, assessment, and remediation processes. The organization operates under a proprietary engagement framework known as the CARE model (Cybersecurity, Assessment, Remediation, and Education). Supported by formal Standard Operating Procedures (SOPs), an elected leadership structure, and functional departments, the Clinic functions year-round as an independent student organization rather than a course-bound initiative. This operational model enables students to join early in their academic careers and progressively refine their expertise over multiple years. Through a corporate-style structure and formalized operational processes, members gain experience in a realistic professional environment and advance into leadership roles. This sustained engagement fosters critical workforce competencies, including professional communication, technical writing, project management, conflict resolution, and performance evaluation. This presentation will outline the structural and operational framework of the UNLV Cyber Clinic, discuss the challenges of managing a student-led cybersecurity clinic, and share proven strategies for sustaining long-term impact.

Keywords: Cybersecurity clinic, student club, cybersecurity career, hands-on training, small business.

Human-centric defense-in-depth framework: Restoring human agency in modern AI-augmented cyber defense strategies

[Presentation]

Petter Lovaas, Niagara University, NY, plovaas@niagara.edu

Perry Benson, Niagara University, NY, pbenson@niagara.edu

Glenn Papp, Niagara University, NY, grp@niagara.edu

Extended Abstract

In recent years, practitioners have witnessed a rush to use artificial intelligence (AI) in defensive and offensive cybersecurity tools and technologies. Companies have deployed AI and machine learning (ML) in threat detection, predictive analytics, and various automated attack response systems. Unfortunately, this environment has fostered novel vulnerabilities such as renewed adversarial attacks on ML models, excessive dependence on algorithmic decision-making, and the atrophy of human analytical skills and expertise within security operations centers (SOCs). Many recent attacks have exploited gaps between security tools and lack of skilled personnel to react to complex adversarial attacks.

AI-driven social engineering attacks have caused major data and financial breaches by exploiting human trust rather than technical flaws. At Arup, deepfake video and voice impersonations of executives led to $25 million in fraudulent transfers, while a compromised PowerSchool account exposed data on 62 million students after AI-assisted credential abuse went undetected. Similar AI-based vishing campaigns later targeted elite universities, underscoring the lack of AI-aware monitoring, training, and verification controls across organizations.

Cybersecurity has long promoted and outlined that humans are one of the weakest links to protecting organizational assets, and this remains true in the age of AI. Accordingly, it is crucial that organizational use of AI helps to improve the expertise of its users through explainable decisions, hypothesis suggestions, and organized analytical methods. To address these gaps, this research presents a reimagination of the defense-in-depth theory through a novel framework that considers how humans can occupy the center, rather than periphery, of AI-augmented security architectures. The proposed Human-Centric Defense-in-Depth framework examine how a human-centric approach can enhance our understanding of the relationship between human cognition, machine intelligence, and resilient security readiness utilizing defense-in-depth as a baseline. The framework that this research proposes builds on the idea of defensive layering using cognitive complementarity, Adaptive Authority Allocation (AAA), and expertise preservation as the layers.

Keywords: Artificial intelligence, human factor, human-centric defense-in-depth framework.

Empowering rural advisors and counselors to guide students toward cybersecurity pathways

[Presentation]

Kristine Christensen, Moraine Valley Community College, IL, christensen@morainevalley.edu

Debasis Bhattacharya, University of Hawaii Maui College, HI, debasisb@hawaii.edu

Extended Abstract

Rural schools face real and ongoing challenges in providing students with meaningful exposure to cybersecurity education and career pathways, contributing to what is often described as "cybersecurity deserts." Many rural schools lack access to a cybersecurity curriculum or qualified faculty and operate with high student-to-counselor ratios, limited staffing and budgets, geographic isolation, and fewer connections to industry and postsecondary partners. Together, these challenges limit counselor capacity and reduce students' awareness of cybersecurity opportunities.

National data show the disproportionate impact of these challenges on rural communities. Research indicates that rural schools have less access to cybersecurity education resources than non-rural institutions, and national workforce reports identify educational gaps as contributing factors to more than 500,000 unfilled cybersecurity positions. Counselors and advisors play an important role in helping students explore academic and career options, yet many report lacking the background or resources to confidently guide students toward cybersecurity pathways.

This presentation introduces Explore Cyber, a toolkit created to support rural school counselors and advisors in guiding students toward cybersecurity pathways with confidence. Designed for users without technical backgrounds, the toolkit can be used in any setting, regardless of whether cybersecurity courses are available. Resources include easy-to-understand descriptions of cybersecurity roles, pathway maps illustrating the progression across educational levels, conversation guides for student meetings, and hands-on activities that help students identify interests, strengths, and next steps. The toolkit also includes an exploratory, research-informed work-role matcher that supports student connections to cybersecurity roles aligned with national workforce frameworks (NICE/DCWF). The toolkit also provides a resource list highlighting extracurricular cybersecurity learning opportunities, capture-the-flag (CTF) competitions, and scholarship programs that provide academic, experiential, and financial access beyond the classroom. By sharing common resources, language, and advising strategies, the toolkit also supports CAE-designated institutions in rural locations by supporting a growing community of practice among rural schools, enabling counselors and advisors to learn from one another and collectively strengthen cybersecurity pathways. This project is funded by NSF award # 2500740.

Keywords: Rural cybersecurity education; career counseling; cybersecurity pathways; workforce alignment.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Large Language Model (LLM)-based Intrusion Detection Systems (IDS): A SOC-focused hybrid architecture and application framework

[Presentation]

Yuksel Celik, University at Albany, State University of New York, NY, ycelik@albany.edu

Sakshi Singh, University at Albany, State University of New York, NY, ssingh29@albany.edu

Sanjay Goel, University at Albany, State University of New York, NY, goel@albany.edu

Extended Abstract

Intrusion Detection Systems (IDSs) face significant challenges in both detection performance and operational feasibility due to high-volume network telemetry, evolving attack behaviors, and severe class imbalance. Although Machine Learning (ML) and Deep Learning (DL) approaches used in current IDS solutions can achieve high accuracy, they often produce high false-positive rates in the imbalanced and dynamic nature of real network traffic, thereby complicating alert management in Security Operations Center (SOC) environments. In addition, these methods provide limited support for critical operational tasks such as contextual alarm interpretation and incident-level correlation across related alerts.

When Large Language Models (LLMs) are considered as an alternative, using them as the primary engine for real-time packet- or flow-level detection is generally unsuitable. This is mainly due to their high computational cost and latency, non-deterministic output behavior, privacy and governance concerns related to sensitive telemetry data, and the need for resilience against manipulations such as prompt injection. Therefore, LLMs are better positioned for analysis-oriented roles such as explanation, correlation, and decision support rather than direct detection.

This study proposes a SOC-focused hybrid IDS framework in which a low-latency ML/DL detection engine is retained for primary detection, while the LLM serves as an intelligent analysis layer. Within this framework, the LLM is used for alarm explanation and standardized reporting, high-level threat interpretation through MITRE ATT&CK tactic and technique mapping, incident narrative generation through multi-alert correlation, and evidence-based triage supported by organizational playbooks and threat intelligence sources through Retrieval-Augmented Generation (RAG). The framework consists of two layers. In the first layer, a low-latency ML/DL detector produces risk scores and candidate labels from session or flow features, and its performance is evaluated using standard metrics such as Precision, Recall, F1-score, PR-AUC, and false-positive rate. In the second layer, the LLM Intelligence Layer is invoked only for high-risk alerts to generate evidence-based explanations, ATT&CK mappings, and recommended response actions. To improve reliability, outputs are constrained to structured JSON, inputs are limited to allowlisted fields, and generated responses are grounded in verifiable sources through RAG. This hybrid approach is expected to improve SOC triage efficiency by reducing the operational burden of false positives and increasing the practical value of IDS outputs.

Keywords: Intrusion detection systems (IDS), deep learning (DL), large language models (LLMs), security operations center (SOC).

Reinventing cybersecurity awareness training using generative AI

[Presentation]

Dmitry Zhdanov, Illinois State University, IL, dzhdano@ilstu.edu

Thomas M. Caldera, OSF Healthcare, IL, thomas.m.caldera@osfhealthcare.org

Mary Elaine Califf, Illinois State University, IL, mecalif@ilstu.edu

Extended Abstract

This project is an industry-academia collaboration between a public research university and a regional hospital system in the United States. The project's innovative core centers around the strategic incorporation of Generative AI (GenAI) within cybersecurity awareness training (CAT) – overall and in healthcare specifically. We are building and evaluating a GenAI-driven training platform that can be used in cybersecurity awareness training programs. We augment the default AI models with healthcare-specific cybersecurity information. Our intent is to make cybersecurity training more effective, interactive, and user-driven.

Our broad conjecture is that the use of a GenAI-enhanced tool for CAT will improve security training outcomes. This conjecture can be further operationalized in the following hypotheses:

1. Use of GenAI CAT will lead to higher engagement with training.
2. Use of GenAI CAT will lead to better retention of training information.
3. Use of GenAI CAT will lead to a better distribution of training in time.
4. Use of GenAI CAT will improve the cybersecurity posture of the organization.

We have developed the GenAI cybersecurity training artifact, CATBOT. We recruited participants from the healthcare system. We collected basic demographic information and performed a baseline security knowledge assessment. After the initial assessment, the study proceeds as follows:

a) Control group – no further actions until knowledge retest.
b) Training group – participates in passive cybersecurity training (watching videos.)
c) Training+AI group – participates in passive training AND interacts with the CATBOT.

A knowledge re-test was performed after two weeks. We have completed a pilot study with 39 participants (out of 300 invited, response rate 13%). The average self-reported computer skills were 4.6/5, cybersecurity skills 3.1/5, and 36% personally experienced a cybersecurity incident. Our results show an improvement of cybersecurity knowledge from 91.5% to 96.2% overall, partially supporting H2. More importantly, there were no dropouts from the Training+AI group, supporting H1. (There were multiple dropouts from the Control and Training groups.) Average engagement duration with the CATBOT was 3.1 questions (with a range of 1-11). 85.7 percent of participants would recommend using the bot. Open-ended feedback was very positive. We are currently working on a bigger study to test hypotheses 3 and 4.

Keywords: Cybersecurity awareness training, generative AI, engagement, learning.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Converting cybersecurity classes to specifications grading

[Presentation]

Mathew J. Heath Van Horn, Embry-Riddle Aeronautical University, AZ, heathvam@erau.edu

Christopher Warner, Embry-Riddle Aeronautical University, AZ, warnerc9@erau.edu

Extended Abstract

The CAE Conference in October 2024 revealed a mismatch between cybersecurity employers and academic programs. Last year, we presented our solution by changing our courses from lecture-based to hands-on learning. We leveraged student involvement, problem-solving labs, and specifications grading. We have been diligently making changes to our courses, collecting data, and publishing the positive results of our work on these issues. However, our audiences were not overly interested in the favorable data; rather, they were interested in how we changed our courses.

Specifications Grading was necessary to accommodate the time needed to develop and host hands-on learning. Specifications Grading, simply put, is the process of recording students' mastery of skills. Specs Grading clearly articulates students' expectations for achieving learning objectives by performing extensive hands-on activities. After 3 weeks of practice, students begin to pre-grade themselves using detailed rubrics focused on the hands-on experience itself. The instructor verifies the student's work using a "met/unmet" rubric, in which all items must be met to receive credit.

On the first day of class, the instructor discusses the rubrics with the students. Any unclear measurements are discussed in class and adjusted accordingly. Furthermore, students can set their final course grade goals. After all, not everyone wants to put in the effort to get "the A". Bloom's Taxonomy is used to categorize the four levels of passing grades. The fundamental D-student must demonstrate recall and understanding of the elements of every learning objective. The average C-student must apply this understanding in hands-on activities. The superior-skilled B-student must demonstrate evaluation and analysis. Finally, the exceptional A-student must innovate by creating a new course component and evaluate their peer's new components.

This presentation will demonstrate how the learning objectives for a network cybersecurity course were converted from traditional grading to a Specifications Grading format. We will discuss how the rubrics are negotiated with students and applied to the course. This is followed by a discussion of how individual coursework components are compiled into grades. We will present the students' opinions on this grading change. Finally, we will ask for volunteers interested in converting their courses and contributing to the body of knowledge by collecting data from their students.

In conclusion, using Specifications Grading has simplified the grading process, which permits more hands-on activities. Students spend more time developing critical thinking, natural inquiry, problem-solving skills, and self-efficacy in an immersive cybersecurity environment, the exact skills sought by employers. This is accomplished without sacrificing normal grade distributions.

Keywords: Cybersecurity, undergraduate, specifications grading, hands-on learning.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Clinic-in-a-Box: AI-generated organizational profiles for scalable cybersecurity experiential learning

[Presentation]

Paul Wagner, University of Arizona, AZ, paulewagner@arizona.edu

Robert Honomichl, University of Arizona, AZ, rjhonomichl@arizona.edu

Extended Abstract

The persistent gap between cybersecurity theory and applied practice continues to limit student readiness for real-world risk management roles, particularly at the secondary and post-secondary levels. Cybersecurity clinics have demonstrated success in addressing this gap by engaging students in authentic, service-oriented risk assessments for under-resourced organizations. Traditional clinic models often require sustained community partnerships, faculty capacity, and operational maturity that are difficult to scale. This extended abstract presents an AI-enhanced Clinic-in-a-Box program that generates realistic organizational profiles to support scalable, scenario-driven experiential cybersecurity learning.

The Clinic-in-a-Box leverages artificial intelligence to dynamically generate synthetic yet contextually accurate organizational profiles, including mission and structure, governance and policy environment, information technology and asset landscape, threat profile, security posture, operational constraints, and risk tolerance. These profiles serve as stand-ins for real clinic clients, enabling students to engage in structured cybersecurity risk assessment activities without the logistical, legal, and ethical challenges of live organizational engagements. AI-generated organizations are aligned with commonly under-resourced sectors such as K-12 education, small businesses, nonprofits, and local government, reflecting environments most frequently served by cybersecurity clinics. Each profile supports iterative exploration of scoping, threat identification, vulnerability analysis, risk prioritization, and mitigation planning consistent with established clinic-based practices.

Clinic-in-a-Box scenarios are deployable through facilitated tabletop exercises, a mobile cyber range, or a web-based interface that supports synchronous, asynchronous, and hybrid delivery. Across modalities, students assume professional roles, collaborate in teams, document assumptions and limitations, and produce client-ready deliverables that mirror authentic cybersecurity consulting and clinic workflows. Learning activities explicitly align with NICE Workforce Framework tasks, knowledge, and skills related to cybersecurity analysis, risk management, and professional communication. This AI-enabled Clinic-in-a-Box approach lowers barriers to experiential cybersecurity education, expands access to clinic-style learning, and supports equitable workforce development. By integrating generative AI, structured risk assessment pedagogy, and flexible delivery models, the program provides a scalable pathway for preparing students to engage confidently and ethically in real-world cybersecurity risk assessment practice while reinforcing the mission of cybersecurity clinics to serve vulnerable organizations.

Keywords: Cybersecurity clinics, experiential learning, ai-generated scenarios, risk assessment pedagogy, workforce development.

SENTINEL: An immersive workforce model for cybersecurity and information technology talent development

[Presentation]

Anthony Lucas, Wake Tech Community College, NC, walucas1@waketech.edu

Sandellyo Kauba, Wake Tech Community College, NC, sakauba@waketech.edu

Extended Abstract

The persistent shortage of entry-level cybersecurity and IT professionals continues to challenge industry and higher education. In response, Wake Technical Community College has pioneered an innovative workforce development model through its National Science Foundation ExLENT-funded program, SENTINEL (Security and Networking Training through Immersive, Novel, and Experiential Learning). Now in its second year and second cohort, SENTINEL is a scalable, non-credit, cohort-based approach designed to accelerate workforce entry while complementing CAE-CD academic pathways. SENTINEL addresses the cybersecurity talent deficit by recruiting learners at the early stages of their information technology career pathways and immersing them in both technical instruction and the professional workforce environment.

Participants engage in applied learning across cybersecurity, advanced networking, and Internet of Things concepts through hands-on labs, simulations, and industry-aligned problem solving. Certification preparation is integrated, but immersion extends beyond classroom instruction into workforce culture and practice. A defining feature of the SENTINEL model is its holistic approach to immersion. An in-house developed mentoring component pairs participants with industry professionals who provide sustained technical guidance, career navigation, and professional socialization throughout the program. This mentoring model is embedded within the curriculum rather than treated as a supplemental activity. Participants also complete structured, paid internships or work-based learning experiences that further immerse them in real-world security and IT operations and strengthen employability skills. Outcomes from the first cohort demonstrate the effectiveness of this approach. The inaugural cohort achieved a 92 percent retention rate, completed 3,525 hours of paid classroom learning and 1,839 hours of paid internships, engaged with 24 industry partners across North Carolina, and was supported by 12 industry mentors. Participants earned 53 industry-recognized certifications and badges and completed four industry site visits. Upon completion, many participants transitioned into Wake Tech's Cybersecurity AAS degree program, secured full-time employment, or had internships extended with employers assuming financial responsibility due to demonstrated value and long-term hiring potential. Early indicators from the second cohort show continued growth across these measures. This presentation will share the SENTINEL model, outcomes, and lessons learned, highlighting how immersive, non-credit workforce programs that integrate mentoring and internships can serve as a replicable strategy for addressing the cybersecurity talent shortage.

Keywords: Cybersecurity workforce development, immersive learning, experiential education, mentoring models, work-based learning, community colleges, non-credit workforce training, internships.

CD Track: Refereed Extended Abstract Proceedings for Mini-Workshops

The digital clean-up challenge

[Mini Workshop]

Suzanne Mello-Stark, Rhode Island College, RI, smellostark@ric.edu

Extended Abstract

The introductory course, Computer FUNdamentals for Cybersecurity, provides students from all academic backgrounds with essential concepts in digital literacy and risk mitigation. By the conclusion of the semester, students possess a practical comprehension of how cybersecurity principles apply to real world scenarios. To ensure these concepts extend beyond the classroom, I developed the Digital Clean-Up Challenge. This structured, three-phase project is mapped to several NSA CAE Knowledge Units, including Personal Cybersecurity (PSY), Social Engineering (SNG), and Cyber Defense (CDP). It leads students through the process of assessing their current routines, identifying specific vulnerabilities, and implementing measurable improvements.

While many students grasp the abstract theory of cyber threats, they often fail to recognize how daily activities, such as using weak passwords or neglecting software updates, expose their personal and professional environments to risk. This workshop demonstrates how a self-adversarial audit shifts theoretical knowledge into practice. The project structure aligns with the Cyber Defense (CDP) Knowledge Unit by requiring students to perform an identification of assets and a personal risk assessment.

The challenge is divided into three primary deliverables. For the first deliverable, students complete a cyber profile and risk assessment. This requires a review of device habits, password management, social media engagement, and phishing recognition to identify gaps in their cybersecurity posture. This phase focuses heavily on the Social Engineering (SNG) Knowledge Unit by highlighting how oversharing and poor digital hygiene facilitate information gathering by adversaries.

For the second deliverable, students develop a personalized action plan detailing specific change, essential tools, and a practical timeline for enhancing their hygiene. For the final deliverable, students implement these changes and reflect on their progress. This includes activating multi-factor authentication, adopting password managers, and adjusting device configurations, all of which are core components of the Personal Cybersecurity (PSY) Knowledge Unit. The Digital Clean-Up Challenge motivates students to convert academic theory into lifelong habits. Participant reviews have been consistently excellent, with students reporting a significantly higher sense of agency and safety in their online lives.

Keywords: Cybersecurity awareness, self-OSINT, attack surface reduction, digital safety, cyber hygiene, social engineering mitigation.

An interactive visualization platform for training cybersecurity analysts to detect subtle concurrency bugs in Rust

[Mini Workshop]

Young Lee, Texas A&M University-San Antonio, TX, young.lee@tamusa.edu

Extended Abstract

Rust is widely adopted for secure software development due to its compile-time memory safety guarantees. However, concurrency-related vulnerabilities persist, often arising from the use of *unsafe* code blocks or, more subtly, through complex interactions between safe abstractions, such as the Cell Reference Smuggling pattern (CVE-2020-36441). Detecting these flaws is critically challenging because standard static analysis tools typically operate at the source code or Mid-level Intermediate Representation (MIR) level, failing to capture the low-level memory access patterns and transformed semantics exposed after compilation. This limitation creates a significant gap in cybersecurity education: how can we effectively train security analysts to locate and comprehend these complicated, compiler-evading concurrency bugs?

This workshop presents a practical, hands-on session centered on a novel methodology that integrates Rust's Low-Level Virtual Machine Intermediate Representation (LLVM IR) IR with Code Property Graphs (CPGs), powered by an interactive graph-based visualization platform. By leveraging the analytical capabilities of the Neo4j graph database, this system enables the detection of race conditions and thread-safety violations—such as the complex "Cell Reference Smuggling" pattern—that is typically concealed within source code. To participate, attendees should bring a laptop to access the platform, which will be provided via a pre-configured containerized environment, ensuring seamless engagement with the workshop's real-world case studies. This approach provides a crucial resource for the CAE-CD community by offering a scalable way to train security analysts to uncover "compiler-evading" bugs, directly addressing a significant gap in current cybersecurity education and strengthening the workforce capability in secure concurrent programming.

By the end of this mini workshop, attendees will possess the theoretical basis and practical skills necessary to uncover and understand memory and thread safety issues in modern software by analyzing post-compilation instruction patterns, thereby strengthening the workforce capability in secure concurrent programming analysis. Participants explore the "Cell Reference Smuggling" pattern, learning how complex interactions between safe abstractions can lead to security flaws despite Rust's compile-time guarantees.

Keywords: Rust, concurrency, LLVM-IR, code property graphs, graph visualization, cybersecurity education, static analysis.

Hands-on AI red teaming: Practical techniques for uncovering vulnerabilities in generative AI systems

[Mini Workshop]

Holly Yuan, University of Wisconsin-Stout, WI, yuanh@uwstout.edu

Extended Abstract

As generative AI systems proliferate across sectors like cybersecurity, healthcare, and education, ensuring their robustness against adversarial attacks has become paramount. This mini workshop introduces participants to AI red teaming, a critical methodology for systematically identifying vulnerabilities in large language models (LLMs) and AI applications. Drawing from established frameworks such as OWASP's Top 10 for LLM Applications and NIST's AI Risk Management Framework, the session emphasizes hands-on exploration of real-world threats, including prompt injection, jailbreaking, and social engineering exploits.

Participants will engage in interactive exercises using open-source tools like Microsoft's PyRIT (Python Risk Identification Tool for Generative AI), enabling automated red teaming workflows for testing models such as OpenAI's GPT-4o. Activities include crafting multi-turn jailbreak prompts (e.g., DAN and Crescendo techniques), encoding attacks via Base64 or ROT13 to bypass guardrails, and evaluating harms across categories like prohibited content, bias, and fabrication. Through guided capture-the-flag (CTF) scenarios, attendees will simulate attacks on LLMs, analyze responses for safety failures, and discuss mitigations informed by recent research on public red teaming models and hierarchical reinforcement learning for automated attacks.

This workshop bridges theoretical insights from cybersecurity and systems theory with practical implementation, fostering multidisciplinary collaboration to address emergent risks in sociotechnical AI ecosystems. For students, it offers unique benefits by building foundational skills in ethical hacking and AI safety, enhancing critical thinking through adversarial simulations, and providing portfolio-ready projects that demonstrate real-world problem-solving—preparing them for careers in AI ethics, cybersecurity, and responsible AI development. By the end, all participants will gain actionable skills to integrate red teaming into AI development lifecycles, promoting safer deployment of generative systems, and walk away with a well-developed Jupyter notebook replicating key exercises, along with a GitHub repository for ongoing replication and extension. No prior coding experience is required, though familiarity with Python is beneficial.

Keywords: AI red teaming, prompt injection, jailbreaking, PyRIT, generative AI safety, OWASP LLM Top 10, vulnerability assessment.

Secure AI-assisted software development

[Mini Workshop]

Aspen Olmsted, Wentworth Institute of Technology, MA, olmsteda@wit.edu

Extended Abstract

The integration of AI pair-programming tools, such as GitHub Copilot and Amazon CodeWhisperer, has fundamentally accelerated software development. However, this speed often comes at the cost of security and architectural integrity, as AI models often generate code that is functional but lacks appropriate bounds, security checks, or alignment with domain-specific constraints. This workshop introduces a rigorous, model-driven approach to secure AI-assisted development that employs UML Stereotypes and the Object Constraint Language (OCL) as defensive guardrails.

Participants will explore a development workflow in which high-level design models serve as the "source of truth" for validating AI-generated output. By applying specialized stereotypes (e.g., `«Secure_Entry»`, `«PII_Data»`, `«Immutable»`), developers can explicitly tag architectural elements with security requirements. These tags are then paired with OCL expressions to define formal, machine-verifiable constraints—such as ensuring a session token is never null or that a transaction amount remains within a specific range.

The workshop focuses on a proactive "Design-First, Prompt-Second" strategy. Participants will learn to: Create custom UML profiles that capture recurring security patterns and risks inherent in AI-generated code; and write precise OCL invariants that can be used to automatically audit or "contract-check" code snippets provided by AI assistants. The methodology will also mitigate LLM Hallucinations by using formal models to detect when an AI's implementation violates the intended system architecture or security policy. Lastly, participants will learn to establish a verification cycle where the developer uses OCL-backed models to "accept" or "reject" AI suggestions based on formal compliance rather than intuition alone.

By the end of this session, attendees will be able to formalize security requirements into actionable models, reducing the risk of "black box" vulnerabilities in AI-written software and ensuring that rapid development does not bypass critical system invariants. The approach presented can be integrated into a software development course or used in a dedicated course focused on secure software development.

Keywords: Secure software development, AI pair programming.

Best practices in the development of a robust program to cultivate the next generation of leaders in cyber and national security

[Mini Workshop]

Ehren Hill, Virginia Tech National Security Institute, VA, ehren@vt.edu

Extended Abstract

The Hume Center for National Security and Technology (Hume Center) was created in direct response to a persistent national challenge: too few qualified U.S. citizens entering federal service and the broader defense and intelligence enterprise to meet mission demand. From its inception, the Hume Center set a clear, student-centered mission to "develop the next generation of national security leaders." Since its founding, the Hume Center has built a robust, scalable talent-development ecosystem that recruits, engages, and prepares students for impactful careers in cyber and national security—supporting Department of War (DoW) and Intelligence Community (IC) missions through research and workforce development.

The Hume Center provides structured opportunities for U.S. citizens to participate in undergraduate and graduate research across interdisciplinary teams that span critical areas such as cybersecurity, resilience, artificial intelligence and machine learning, and autonomy; engage with government and industry partners on applied problems; and access professional development through seminars, workshops, career fairs, and internships. These offerings are intentionally designed to reinforce one another: seminars and career development activities drive early interest; research and project-based learning build technical depth; and internships and workforce development programs create pathways to meaningful and impactful careers in national security. Together, these components form a coherent pipeline that moves beyond ad hoc engagement to a sustained model of recruitment, development, and placement.

A distinguishing feature of this model is its focus on growth and durability, which informed program design choices including repeatable engagement structures that support scale and partner-driven opportunities that keep student work anchored to real mission needs. This approach has enabled the Hume Center to engage approximately 1,000 students annually—while preserving quality, mentorship, and mission relevance through advising, structured project experiences, and partner participation.

For the CAE Symposium, this workshop will share the Hume Center model and engage participants in a discussion about replicable best practices. The workshop will explore strategies for launching and maturing a student-focused research center, building a recruitment-to-career pipeline, aligning programs to stakeholder requirements, and sustaining growth through partnerships. The objective is to equip other institutions with practical approaches for developing their own national security workforce programs—strengthening the national talent base in cyber and national security to meet evolving DoW and IC mission needs.

Keywords: Cybersecurity, national security, workforce development, undergraduate research, graduate research, experiential learning.

From capture-the-flag to careers: Student clubs as launchpads for cybersecurity success

[Mini Workshop]

Suzanna Schmeelk, St. John's University, NY, schmeels@stjohns.edu

Denise Dragos, St. John's University, NY, dragosd@stjohns.edu

Kutab Thakur, St. John's University, NY, thakurk@stjohns.edu

Joan E. Debello, St. John's University, NY, debelloj@stjohns.edu

Erald Troja, St. John's University, NY, trojae@stjohns.edu

Extended Abstract

Student-led cybersecurity clubs mentored by cyber faculty and industry-leads—especially those that compete in Capture-the-Flag (CTF) events—are uniquely positioned to translate classroom theory into workforce-ready capability. This workshop demonstrates a programmatic approach for aligning club activities with the NIST NICE Workforce Framework for Cybersecurity (SP 800-181 Rev. 1) and the National Centers of Academic Excellence in Cybersecurity (NCAE-C / CAE-CD) Knowledge Units and KSA expectations, so faculty and club leaders can intentionally develop talent pipelines from first meeting to first job offer.

The mini-workshop deliverables include: (1) insights into building a customized matrix linking selected CTF categories to at least 3 NICE Work Roles and 5–8 Competency Areas, plus the corresponding CAE-CD KUs each activity supports; (2) rubrics that tag artifacts to NICE TKS statements and CAE KU outcomes, suitable for use in e-portfolios and CAE evidence binders; (3) template(s) aligning CTFs, guest talks, and scrimmages to NICE competencies with checkpoints timed to CAE documentation needs. Linking CTF practice to NICE and CAE creates a transparent, portable skills signal across academia and industry, strengthens CAE designation dossiers, and shortens time-to-contribution for early-career hires through competency-focused preparation.

Keywords: Capture-the-flag, cyber security workforce and career development, NIST NICE Competency, CAE KSAs.

References:

Cybersecurity and Infrastructure Security Agency. (n.d.). NICE Workforce Framework for Cybersecurity (NICE Framework) [Interactive tool]. *National Initiative for Cybersecurity Careers and Studies (NICCS)*. https://niccs.cisa.gov/tools/nice-framework

National Centers of Academic Excellence in Cyber Defense. (2024, December 9). CAE-CD 2024 Knowledge Units [PDF]. Department of Defense. https://dl.dod.cyber.mil/wp-content/uploads/cae/pdf/unclass-cae-cd_ku.pdf

Petersen, R., Santos, D., Smith, M. C., Wetzel, K. A., & Witte, G. (2020, November). Workforce Framework for Cybersecurity (NICE Framework) (NIST Special Publication 800-181 Rev. 1). *National Institute of Standards and Technology.* https://doi.org/10.6028/NIST.SP.800-181r1

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Development of a self-hosted cyber range to teach and assess cybersecurity competencies

[Mini Workshop]

Douglas Rausch, Bellevue University, NE, drausch@bellevue.edu

Eric Jackson, Bellevue University, NE, ericjackson@bellevue.edu

Extended Abstract

The increasing demand for job-ready cybersecurity professionals underscores the critical importance of practical skills acquisition within college curricula. Despite familiarity with theoretical concepts, many students lack the confidence and hands-on experience necessary to excel in real-world positions, while employers emphasize the need for graduates who can contribute immediately upon entering the workforce. Existing training solutions often present challenges such as prohibitive costs, limited flexibility, and insufficient fit for asynchronous and online learning modalities.

This workshop introduces the development of a self-hosted cyber range for higher education institutions, designed to address these challenges through an accessible, customizable, and competency-based educational platform. The system features instructor-developed, goal-oriented labs accessible via web browser, and supplements student learning with integrated AI tutoring for adaptive support. Built on a Django web framework with AI capabilities provided by Ollama, and virtualized VMware backend, the platform enables scalable, formative, and summative assessments directly aligned with the Department of Defense Cyber Workforce Framework (DCWF).

Early student trial surveys indicate that the ease of range use not only fosters a supportive environment where students feel comfortable practicing and receiving valuable, constructive feedback, but also builds strong confidence in their ongoing task progression. Initial labs have been developed for Introduction to Operating Systems (Configuration of a Linux OS), Incident Detection and Response (Introduction to Ghidra and YARA rules) and Penetration Testing (Pentesting of a small network). Future work will focus on expanding the breadth of task-based labs, providing additional AI tutor personas and progressing from pilot testing to widespread deployment. Deployment of the range framework on other than VMware based infrastructure should be possible with minimal modifications with a test deployment to a Proxmox virtualized environment currently underway. By empowering both instructors and students, this tailored cyber range provides a flexible and cost-effective solution for building industry-relevant cybersecurity skills. This approach positions institutions to bridge the gap between academic preparation and workforce readiness in the dynamic field of cybersecurity.

Keywords: Cyber range, competencies, Defense Cybersecurity Workforce Framework, workforce development, cybersecurity skills, virtualization, self-hosting.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Embedding AI-enabled workflows into cybersecurity curriculum: A practical integration model

[Mini Workshop]

Elliott S. Lynn, American Public University System, WV, Elliott.lynn@mycampus.apus.edu

Extended Abstract

As artificial intelligence tools become increasingly embedded in cybersecurity practice, academic programs face growing pressure to integrate these technologies into curriculum in ways that are pedagogically sound, ethically responsible, and aligned with workforce expectations. While many institutions recognize the importance of AI literacy, faculty often encounter uncertainty regarding how to incorporate AI-enabled workflows without undermining learning objectives, assessment integrity, or foundational cybersecurity competencies. This challenge is particularly pronounced within cybersecurity education, where technical rigor and professional accountability are central.

This mini workshop presents a practical, skills-focused model for embedding AI-enabled workflows into cybersecurity curriculum as an integrated component of learning rather than as a standalone topic. The session emphasizes how AI tools can support core cybersecurity activities such as analysis, documentation, decision support, and risk evaluation while reinforcing critical thinking and domain expertise. Rather than teaching students to rely on automated outputs, the model prioritizes guided interaction with AI tools, structured evaluation of results, and explicit consideration of ethical, legal, and policy constraints.

The workshop introduces a curriculum integration framework that aligns AI-enabled tasks with existing course objectives, assessments, and learning outcomes commonly found in CAE-CD programs. Attention is given to instructional design strategies that promote transparency, accountability, and academic integrity while allowing students to demonstrate applied AI fluency within cybersecurity contexts. The discussion also addresses how AI-enabled workflows can be incorporated across technical, governance, and management-oriented courses without displacing essential cybersecurity content.

By focusing on implementation strategies rather than tool promotion, this workshop contributes to the CAE-CD community's efforts to modernize cybersecurity education while preserving rigor and consistency. Attendees will gain a repeatable approach for integrating AI-enabled workflows into cybersecurity curriculum that supports skill development, ethical awareness, and workforce readiness across diverse program structures.

Keywords: Cybersecurity curriculum integration, artificial intelligence workflows, instructional design, academic integrity, cybersecurity education.

Designing robust systems: Secure and defensive programming in Rust

[Mini Workshop]

Christian Servin, El Paso Community College, TX, cservin1@epcc.edu

Extended Abstract

As software systems increasingly underpin critical infrastructure, the need for programming practices that prioritize security, reliability, and resilience has never been greater. Rust has emerged as a compelling language for both cybersecurity education and general computing due to its strong guarantees around memory safety, concurrency, and correctness -- without sacrificing performance. By eliminating entire classes of vulnerabilities such as *buffer overflows*, *use-after-free errors*, and *data races* at compile time, Rust directly addresses many of the root causes of modern software exploits.

This workshop presents a series of hands-on instructional modules that teach secure programming and cybersecurity best practices using Rust, explicitly aligned with CAE Knowledge Areas including Secure Software Development, Programming, Systems Security, and Vulnerability Analysis (Cybersecurity, 2024). The modules emphasize adversarial and defensive thinking through practical examples such as safe handling of untrusted input, secure memory management without garbage collection, race-condition-free concurrent programming, and the design of robust system components.

Participants will explore how Rust's ownership model, strong type system, and compiler-enforced guarantees naturally reinforce secure coding behaviors and reduce common implementation flaws. The integrating Rust into foundational and advanced computing curricula, this workshop demonstrates how secure software development can be embedded as a first-class concern aligned with CAE outcomes—preparing students and practitioners to design, implement, and reason about software that is secure by design across modern computing environments.

Keywords: Rust programming, secure and reliable programming.

References:

National Centers of Academic Excellence in Cybersecurity. (2024). *Cyber Defense (CAE-CD) Knowledge Units (KUs) (Version 2024.06)*. Department of Defense. https://dl.dod.cyber.mil/wp-content/uploads/cae/pdf/unclass-cae-cd_ku.pdf

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Towards a unified fine-grained access control model for research computing infrastructure

[Mini Workshop]

Indrakshi Ray, Colorado State University, CO, indrakshi.ray@colostate.edu

Mahmoud Abdelgawad, Colorado State University, CO, m.abdelgawad@colostate.edu

Abhimanyu Chawla, Colorado State University, CO, abhimanyu.chawla@colostate.edu

Extended Abstract

Modern scientific research is often data-driven and collaborative in nature, involving multiple institutions, possibly spanning different countries. However, data, algorithms, and experimental results are often sensitive and subject to legal and organizational disclosure restrictions. Collaborative projects require sharing data and resources, but this must be carefully managed to protect users and organizations, especially given current infrastructures. Traditional access control models used in research collaboration environment often lead to overly broad access, over-provisioning, slow revocation, and increased risks of misuse, leakage and violating legal requirements. Most organizations use a combination of identity-based access control (for accountability) and role-based access control (for simpler policy management). We argue that a uniform access control model would simplify policy management, reveal conflicts and redundancies, and potentially enable partial automation of policy enforcement. We propose an attribute-based access control model, NGAC++, derived from the NIST Next Generation Access Control, to secure the Research Computing Infrastructure (RCI). Access is determined by subject attributes, resource characteristics, and environmental conditions. The model supports fine-grained, cross-organizational access control and is modular, enabling consistent use across databases, operating systems, and research applications. We explain how to derive the access control model from security requirements, analyze it for well-formedness, conflicts, and redundancies (Alqurashi, Abdelgawad, Shirazi, 2024). We demonstrate how NGAC++ can be used to protect RCI (Chawla, Abdelgawad, Ray, 2025). We are currently working on implementing the framework for an example RCI and demonstrating how system logs can be used to verify enforcement and fine-tune the NGAC++ model.

Keywords: Next generation access control, research computing infrastructure.

References:

Alqurashi, S., Ray, I., Abdelgawad, M., Shirazi, H., (2024, Nov). SR2ACM: A methodical approach for translating security requirements to access control model. *Proceedings of the 6th IEEE International Conference on Trust, Privacy and Security in Intelligent Systems, and Applications*.

Chawla, A., Abdelgawad, M., Ray, I., (2025, Nov.) Access control policies specification and analysis for multi-institutional collaborative projects. *Proceedings of the 11th IEEE International Conference on Collaboration and Internet Computing*.

Toward an academic rigor analysis for cybersecurity education

[Mini Workshop]

Ram Dantu, University of North Texas, TX, ram.dantu@unt.edu

Cihan Tunc, University of North Texas, TX, cihan.tunc@unt.edu

Fatima Shibli, University of North Texas, TX, fatima.shibli@unt.edu

Extended Abstract

Cybersecurity education and workforce readiness are major concerns as cyberattacks become more frequent and complex, affecting not only individuals or companies but also critical infrastructures and government agencies. However, cybersecurity training content can differ drastically across institutions, professors, semesters, and student populations. This discrepancy creates unknown levels of preparedness in the workforce and can lead to missing skills expected of students, costing both sides time and money in retraining and productivity loss.

Even though rigor in cybersecurity education has been discussed at the policy, accreditation, and course design levels, it has been addressed only rarely. This is especially the concern for rigor quantification. Furthermore, the existing methods typically gauge student performance/outcome rather than assessments (assignments/lab work/projects). For this purpose, workshop will discuss how rigor analysis can be defined and measured in cybersecurity education and demonstrate, in an interactive setting, how large language models (LLMs) can be used as a scoring instrument, etc. Specifically, the workshop will focus on two main components. First, attendees will interactively evaluate the rigor of cybersecurity courses across semesters and instructors, discussing how rigor may vary and identifying potential methods to improve its evaluation. Second, we will present our approach to rigor analysis that benefits from LLM as a calibrated measurement tool to assign Knowledge, Skill, and Ability (KSA) scores based on Bloom's taxonomy using instructor-created artifacts such as homework, labs, and projects. In this framework, the LLM acts as a calibrated measurement instrument that evaluates the cognitive and technical demands embedded in each assessment using a fixed prompt-engineered rubric.

When applied across multiple advanced cybersecurity and computing courses, *we observe that* the resulting learning trajectories exhibit consistent structural patterns, including monotonic growth, early acceleration, diminishing returns, and late-stage saturation. This supports the intuitive idea that, in most cybersecurity courses, students learn new material early in the semester and apply it by the end of the semester. Our solution enables these patterns parameterized for homework-intensive, lab-intensive, project-intensive, and hybrid courses, enabling cross-course comparison, curriculum auditing, and evidence-based reflection for cyber defense education, program improvement, and accreditation. In summary, the proposed workshop will discuss the importance of rigor, methods for rigor analysis, and our solution to rigor analysis using an LLM as a measurement tool.

Keywords: Academic rigor, cybersecurity education, assessment analytics, large language models, Bloom's taxonomy, knowledge-skill-ability modeling.

iEXAM: AI-Driven cybersecurity with explainability

[Mini Workshop]

Indrakshi Ray, Colorado State University, CO, indrakshi.ray@colostate.edu

Rakesh Podder, Colorado State University, CO, rakesh.podder@colostate.edu

Sarath Sreedharan, Colorado State University, CO, sarath.sreedharan@colostate.edu

Indrajit Ray, Colorado State University, CO, indrajit.ray@colostate.edu

Shadaab Kawnain Bashir, Colorado State University, CO, shadaab.bashir@colostate.edu

Extended Abstract

Graph-based frameworks are often used in network hardening to help understand how a network can be attacked and how best defenses can be deployed. Tools built around such frameworks can potentially help students to evaluate defensive strategies and techniques, understand tradeoffs in implementing those and refine their skill. However, most works fail to incorporate network connectivity information or allow reasoning in the absence of complete information. These are important because a defense strategy may result in a loss in connectivity or bring up new issues thus rendering the strategy unacceptable. Existing frameworks also do not allow experiments with different what-if analysis scenarios, considering various configuration and attacker motives and deliver suggestions to improve security postures.

Towards this end, we present i-EXAM, a tool that supports interactive security profiling and what-if analysis for enterprise networks. We use our earlier Attack Connectivity Graph (ACG) framework (Podder et al., 2025) for representing attack and network connectivity information. We use AI planning techniques to analyze various what-if scenarios, experimenting with different configurations and attacker objectives. We use state-of-the-art large language model (LLM) to explain these results to the user. i-EXAM automatically constructs ACGs from host and network data and vulnerability intelligence, compiles them into Planning Domain Definition Language (PDDL) problems, and uses planners to compute attack paths, evaluate security postures, and synthesize diverse hardening options while preserving required services. It explains recommendations via model restriction and LLM generated, actionable rationales. For cyber defense education, this ability to perform interactive what-if analysis is a very useful feature since it allows the student to easily verify and validate solutions and understand tradeoffs in implementing defenses.

Acknowledgment: Partially supported by the U.S. ONR under award #N000142612041.

Keywords: Network Hardening, Attack Graph Analysis, AI Planning, Explainability

References:

Podder, R., Caglar, T., Bashir, S. K., Sreedharan, S., Ray, I., & Ray, I. (2025, July). SPEAR: security posture evaluation using AI planner-reasoning on attack-connectivity hypergraphs. *Proceedings of the 30th ACM Symposium on Access Control Models and Technologies* (pp. 62-73).

From cyber attacks to system recovery: A resiliency-centered view of cyber-physical systems security

[Mini Workshop]

Indrakshi Ray, Colorado State University, CO, indrakshi.ray@colostate.edu

Shadaab Kawnain Bashir, Colorado State University, CO, shadaab.bashir@colostate.edu

Rakesh Podder, Colorado State University, CO, rakesh.podder@colostate.edu

Sarath Sreedharan, Colorado State University, CO, sarath.sreedharan@colostate.edu

Indrajit Ray, Colorado State University, CO, indrajit.ray@colostate.edu

Extended Abstract

Digitalization of industrial control systems (ICS) has significantly improved automation and operational efficiency but has increased the attack surface in cyber-physical systems (CPS) that can potentially affect safety. Safety and cybersecurity analyses are carried out separately in CPS. This prevents operational technology (OT) operators/learners to have a thorough grasp of the impact of cyber threats on the OT. We present a framework which facilitates methodical analysis of how cyber-attacks impact CPS behavior, safety and recovery. The framework models the causal relationships between cyber vulnerabilities, system configurations, and physical failures within a unified transition-system representation. It enables users to reason about why specific cyber-attacks lead to disproportionate operational or safety consequences and how resiliency can be preserved or restored.

In the framework, cyber-attacks, misconfigurations, and component failures are represented as actions of an AI planning model with explicit preconditions and effects. Adversarial behavior is integrated with failure, enabling a systematic analysis of how attack to failure cascades. The resulting transition system is encoded in a domain-specific AI-planning language. It utilizes off-the-shelf planning engines to explore alternative system evolutions under attack. Planning-based reasoning is instrumental in crucial cyber resiliency analysis, including identifying critical assets whose compromise significantly diminishes resiliency, conducting what-if analyses under varying attacker objectives or incomplete information, and generating diverse mitigation and recovery strategies. These strategies are accompanied by natural-language explanations that elucidate the causal relationships between cyber events and physical outcomes. Evaluations on a representative ICS configuration (a natural gas flare system, see Bashir et al., 2024.) demonstrate that this approach effectively captures complex cyber-physical dependencies while maintaining scalability.

Acknowledgment: Partially supported by the U.S. ONR under award # N000142612041.

Keywords: Cyber physical system, cyber-attacks, AI planning, resiliency analysis.

References:

Bashir, S. K., Podder, R., Sreedharan, S., Ray, I., & Ray, I. (2024). Resiliency graphs: Modelling the interplay between cyber-attacks and system failures through AI planning. *Proceedings of the IEEE 6th International Conference on Trust, Privacy and Security in Intelligent Systems and Applications (TPS-ISA)* (pp. 292–302).

Competency in credentials: Calculating proficiency in certificates using large language models

[Mini Workshop]

Alexis Blackwell, University of North Texas, TX, alexisblackwell@my.unt.edu

Ram Dantu, University of North Texas, TX, ram.dantu@unt.edu

Extended Abstract

Industry credentials are an important metric for assessing an individuals' mastery of skills, knowledge, and competency in a certain field. Certificates, certifications, and licenses are all different types of credentials that can be used as proof of skill mastery. However, there are a vast number of credentials available, each covering various topics. There is also little standardization, as some credentials may offer a basic overview of a topic whereas others may expect an individual to learn advanced theory and skills.

Current methods for evaluating proficiency rely on either manual assessments or using general learning taxonomies to estimate learning levels. These learning levels can then be mapped to proficiency levels. The proficiency levels are Basic, Intermediate, and Advanced. Basic is considered a fundamental understanding of a skill or topic, Intermediate is a moderate amount of understanding, and Advanced is a deep understanding of the material. However, the use of these general learning taxonomies is not ideal as domain specific verbs may be missing or there may be missing context around a verb (Adeleye et al., 2023). This can lead to estimating both the learning level and proficiency level. Previous research on this topic focused on adapting the general learning taxonomies to be domain specific by adding more verbs to the existing taxonomy (Adeleye et al., 2023).

A proposed solution is to use a Large Language Model (LLM) with Retrieval-Augmented Generation (RAG) and Knowledge Graphs. This methodology uses Hybrid Text Summarization to extract and summarize each topic in the certificate. RAG and Knowledge Graphs are used to provide domain specific context for the LLM to evaluate Cognitive and Psychomotor learning levels with topic complexity. These Knowledge Graphs were generated using an LLM to extract triples from a Computer Science Curricula Recommendations corpus. Proficiency is calculated using the learning levels and topic complexity. The dataset consists of 43 labeled Cybersecurity certificates. A random selection resulted in assigning 92% of those certificates with multiple levels of proficiency while 50% of those certificates had been manually labeled as having a single level of proficiency. Overall, this process shows how LLMs can be used to analyze individual topics within a credential to create a higher resolution of a credential's proficiency level.

Keywords: Large Language Models (LLMs), technical proficiency, knowledge graphs, learning taxonomy, Retrieval-Augmented Generation (RAG).

References:

Bamkole, A., Geissler, M., Koumadi, K., Servin, C., Tang, C., & Tucker, C. S. (2023). Bloom's for computing: Enhancing Bloom's revised taxonomy with verbs for computing disciplines. https://doi.org/10.1145/3587276

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Cyber range proficiency and rigor assessment: An automated framework with LLM-enhanced psychometric validation

[Mini Workshop]

Elijah Goodrich, University of North Texas, TX, elijahgoodrich@my.unt.edu

Chad Mello, United States Air Force Academy, CO, chad.mello@afacademy.af.edu

Ram Dantu, University of North Texas, TX, ram.dantu@unt.edu

Extended Abstract

Cyber ranges are essential training environments for developing cybersecurity skills, yet they often lack automated, psychometrically grounded measurement of scenario difficulty and multi-dimensional participant proficiency. This limitation restricts the rigor, comparability, and actionable feedback that CAE-CD institutions can provide, making it difficult to validate learning outcomes and provide evidence-based skill development feedback.

This workshop presents an end-to-end automated assessment framework that addresses these challenges through LLM-enhanced psychometric validation. The framework computes scenario rigor R(s) as a weighted composite of five components: infrastructure fidelity, attack complexity, learning coverage (ATT&CK-to-DCWF KSAT mapping), data-driven difficulty, and LLM semantic complexity. Participant proficiency is computed as $P\_i = \alpha \cdot P\_Q + \beta \cdot P\_L$, blending quantitative telemetry with structured LLM evaluation. Psychometric validity is ensured through 2PL IRT model calibration using expectation-maximization with EAP estimation (Dempster et al., 1977; Lord & Novick, 1968). In red-vs-blue experiments (8 red, 8 blue agents; 100 iterations completed), the red team achieved an average composite score of 0.481 (min–max 0.417–0.517) and the blue team 0.304 (0.268–0.390), with red score improving by +0.009 and blue by −0.026 when comparing the first 10 to the last 10 iterations. At the final iteration, detection coverage of executed ATT&CK techniques was 84.6% (11 of 13 techniques detected); attack sequences spanned 30 unique ATT&CK technique IDs mapped to 30 DCWF KSATs. Joint payoff was 0.383 and cooperativity 0.69. Scenario difficulty was held constant at 0.60 in this run.

This workshop provides hands-on experience with scenario rigor calculation, multi-modal proficiency assessment, IRT calibration, and LLM-enhanced evaluation. Participants will learn to implement automated assessment in their cyber range environments and generate actionable feedback for skill development.

Keywords: cyber range, proficiency assessment, item response theory, psychometric validation, LLM evaluation, automated assessment, DCWF mapping

References:

Bianchi, F., Bassetti, E., & Spognardi, A. (2023). Scalable and automated evaluation of Blue Team cyber posture in cyber ranges. arXiv. https://doi.org/10.48550/arXiv.2312.17221

Lillemets, P., Jawad, N. B., Kashi, J., Sabah, A., & Dragoni, N. (2025). A systematic review of cyber range taxonomies: Trends, gaps, and a proposed taxonomy. *Future Internet, 17*(6), 259. https://doi.org/10.3390/fi17060259

CD Track: Refereed Extended Abstract Proceedings for Lightning Talks

Transforming cybersecurity education with artificial intelligence

[Lightning Talk]

Laura Hill, College of Western Idaho, ID, laurahill1@cwi.edu

Scott Didriksen, College of Western Idaho, ID, scottdidriksen@cwi.edu

Extended Abstract

Artificial Intelligence (AI) is revolutionizing cybersecurity education in colleges and universities, reshaping how students are prepared for digital defense careers. Academic programs increasingly leverage AI technologies to modernize curriculum delivery, enhance experiential learning, and align student competencies with workforce expectations. AI's impact is evident in campus-based Security Operations Centers (SOCs), where AI-driven tools provide hands-on experience in threat detection and incident response. Institutions like Oregon State University and Auburn University have reported that AI accelerates student onboarding, improves detection rule creation, and empowers students to contribute to live SOC operations earlier.

Microsoft and Splunk did industry analysis and validated the need for SOC education noting it reduced institutional security costs. AI is also transforming instructional design through adaptive learning platforms, automated labs, and predictive analytics that simulate real-world attack scenarios. These tools support differentiated instruction, tailoring content to individual needs while maintaining rigor.

Challenges remain, including faculty expertise gaps, resource limitations, and the need for strategic partnerships to align curricula with industry standards. To address these barriers, the College of Western Idaho (CWI) is adopting a multi-dimensional framework combining AI-enhanced pedagogy, faculty development, and collaborative SOC models. This holistic approach fosters student competency in roles like cybersecurity analyst and AI security engineer, positioning institutions as hubs for workforce development. This includes curriculum modernization through AI-driven tools, expanded experiential learning via SOCs, faculty development initiatives, and strengthened industry partnerships. AI is a strategic imperative for cybersecurity education. By embedding intelligent systems into instruction and operations, higher education institutions are cultivating cyber professionals prepared to defend digital infrastructure and innovate in a rapidly evolving field, ensuring graduates possess the skills to address complex cybersecurity challenges.

Keywords: Artificial intelligence, cybersecurity education, experiential learning, SOCs.

Cyber-physical systems in cybersecurity education: The Capture the Smart TAG Competition (CSTC) case study

[Lightning Talk]

Stan Mierzwa, Kean University, NJ, smierzwa@kean.edu

Extended Abstract

Cyber-Physical Systems (CPS) are increasingly central to modern cybersecurity education, bridging the gap between theoretical knowledge and real-world application. The Capture the Smart TAG Competition (CSTC), hosted by Kean University's Center for Cybersecurity in partnership with Berkeley Varitronics Systems, exemplifies this integration by immersing students in a hands-on challenge that combines Bluetooth Low Energy (BLE) technology, physical security awareness, and digital threat detection. The CSTC addresses a gap in cybersecurity exercises and engages student teams in a timed search for hidden smart tags—such as Apple AirTags and Tile Trackers—using the BlueSleuth-Lite BLE Tag Detector throughout a campus. The student competition unfolds in multiple rounds, with teams of 3–4 students navigating designated campus areas to locate discreetly placed tags.

A golden tag offers bonus points, and the format includes single-elimination and timed challenges. Participants receive training on BLE technology, collaborate on strategy, and compete under the supervision of cybersecurity professionals from the CPS professional industry and non-profits, such as ISC2, FBI InfraGard, and ISACA. This event not only cultivates technical and soft skills—such as teamwork, problem-solving, and situational awareness—but also raises awareness of the dual-use nature of smart tracking devices. CSTC demonstrates how CPS-based learning can enhance cybersecurity curricula by simulating real-world threats in a controlled, engaging environment. The lightning talk of this case study, aligned with workforce readiness and innovative cybersecurity education, will outline the CSTC's structure, educational impact, and replicable format, offering a blueprint for integrating CPS into cybersecurity programs to better prepare students for evolving threats in both digital and physical domains.

Keywords: Cyber-Physical Systems (CPS), Smart TAGS, Bluetooth Low Energy (BLE), Competitions, power skills.

Closing the readiness gap: Preparing cybersecurity students to stand out in a crowded job market

[Lightning Talk]

Jason Mitchell, Lansing Community College, MI, mitch24@lcc.edu

Extended Abstract

Community colleges play a critical role in the national cybersecurity ecosystem, serving as a primary entry point into the profession for diverse learners, including career changers, first-generation students, and working adults. Over the past decade, these institutions have expanded program capacity and increased the number of technically prepared graduates entering the workforce. As a result, the entry-level labor market has evolved: while demand for experienced practitioners remains strong, employers increasingly report that new graduates often possess similar technical skills but vary widely in their readiness to operate effectively in professional environments. This shift reframes the central challenge from a shortage of talent to a gap in professional readiness.

This lightning talk introduces a practice-oriented framework that positions workforce readiness as the next evolution in cybersecurity education, particularly within community colleges where programs are closely aligned with regional workforce needs. The novelty lies in reframing employability not as a function of additional certifications or deeper technical specialization, but as the intentional cultivation of four core competencies—Capability, Confidence, Communication, and Curiosity—that complement technical training and serve as differentiators in hiring decisions.

The presentation highlights how these competencies are operationalized within a CAE-CD-aligned curriculum through portfolio-producing assignments, practicum-style experiences, communication-focused deliverables, and reflective learning opportunities. Curiosity is encouraged through exploratory labs and student-driven investigations that promote continuous learning habits. Early program observations, including employer feedback, student placement outcomes, and increased student engagement, suggest that integrating these elements strengthens graduate readiness. This session presents a scalable model for CAE-designated programs to better align educational outcomes with evolving workforce expectations and prepare graduates to contribute effectively from day one.

Keywords: Cybersecurity education, workforce readiness, community colleges, employability skills, student development, readiness gap.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Establishing a security operations center

[Lightning Talk]

Rob Greenberg, Sam Houston State University, TX, rjg046@shsu.edu

Extended Abstract

Institutions of higher education are increasingly subject to cyberattacks, joining government agencies, private enterprises, and individuals as frequent targets. Universities are particularly vulnerable due to their open networks, diverse user populations, and the sensitive data they manage. Despite these risks, many universities possess substantial internal expertise in cybersecurity through their computer science departments and student populations.

Recognizing this alignment of institutional need and academic capability, we chose to create a Security Operations Center (SOC). A SOC provides real-time network security monitoring, analyzes suspicious network traffic, and initiates responses when security incidents are detected. These responses range from issuing alerts and guidance to coordinating formal incident responses involving malware infections, unauthorized access attempts, or ransomware attacks.

The development of a SOC offers substantial benefits to multiple stakeholders, including the institution, students, the academic department, and the broader community. This paper presents a case study of the creation of a university-based SOC and outlines practical, replicable steps for establishing similar centers at other institutions of higher education. Creating a SOC should be a formal project requiring defined leadership, a dedicated project team, and a comprehensive project plan. Such a plan must identify critical challenges and outline strategies for addressing them. Initial phases should focus on identifying key institutional partners and sponsors, securing their support, and aligning the SOC's mission with institutional priorities.

Financial considerations must be addressed early in the process. Operating a SOC typically requires paid student employees, and costs increase significantly when year-round coverage is desired. Both one-time and recurring expenses, including staffing, software, hardware, and training, must be projected, and sustainable funding sources must be secured. Staffing represents another major challenge. Suitable students must be recruited, hired, and trained to operate the SOC effectively. In addition, ongoing supervision and mentorship are essential to ensure both operational success and educational value. Confidentiality concerns must also be addressed given the nature of the data (FERPA, etc.) flowing through the SOC. Like all university employees exposed to such data, they can be addressed through a combination of training and signed agreements.

Finally, successful implementation depends on the ability to clearly articulate the value proposition of the SOC to institutional stakeholders. This includes demonstrating how the SOC enhances cybersecurity posture, supports the campus community, contributes to regional workforce development, integrates real-world experiences into academic curricula, and provides a foundation for future expansion, such as offering services to external organizations.

Keywords: SOC, security, workforce development, real-world training.

Using AI tools to build cybersecurity curriculum and create instructional tools

[Lightning Talk]

Kasia Taylor, Anne Arundel Community College, MD, ktaylor2@aacc.edu

Mary McDonald, Anne Arundel Community College, MD, memcdonald2@aacc.edu

Extended Abstract

The foundational computing curriculum in our cybersecurity programs required a broader range of cybersecurity concepts to strengthen students' preparedness for higher-level coursework. To support this process, we used AI tools CoPilot, ChatGPT and Notebook LM and concluded that the same technologies employed to enhance curriculum relevance could also be leveraged for the development of instructional resources. Our plan is to expand the scope and depth of cybersecurity topics integrated into the program's foundational computing course, while creating additional instructional resources to facilitate applied learning and skill acquisition using chatbots, podcasts and video overviews.

The initiative encompasses four key components. First, curriculum development will leverage artificial intelligence to integrate cybersecurity relevance across all topics within the foundational computer technology course, ensuring a comprehensive and security-conscious framework. Second, instructional support will employ AI-driven solutions to design pedagogical tools, including interactive chatbots, multimedia overviews, and structured study guides, aimed at enhancing learner engagement and comprehension. Third, data analysis will involve a comparative evaluation of course objectives between sections adopting the revised structure and those following the original design, providing empirical evidence of effectiveness. Finally, a survey was conducted to gather insights from other colleges and universities regarding their current utilization of AI or future-plans to incorporate AI technologies into cybersecurity education. The results will be shared at the symposium.

The lighting talk will examine the application of specific artificial intelligence tools in shaping both curriculum design and instructional resources within the program. It will highlight the AI technologies that demonstrated the greatest impact on pedagogical effectiveness and curriculum integration. Additionally, the session will present a data-driven analysis of learning objectives associated with the revised course structure, followed by a discussion of survey findings that capture trends in the adoption or planned implementation of AI within cybersecurity education.

Keywords: Foundation, cybersecurity, curriculum, design, AI, instruction, resource.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Guided AI for CTF-based cybersecurity education: A roadmap for reasoning, efficiency, and integrity

[Lightning Talk]

Deep Ramanayake, Xavier University, OH, ramanayaked@xavier.edu

Extended Abstract

Cybersecurity students are increasingly asking AI tools to interpret logs, debug exploits, and solve capture-the-flag (CTF) challenges. These tools can reduce friction and broaden access to expert-like help, yet unguided use introduces risks such as over-reliance, hallucinated steps, and shallow reasoning that does not transfer to novel problems. At the same time, blanket "no AI" policies are becoming unrealistic as AI assistance spreads across academic and professional cybersecurity practice. This lightning talk presents a narrative synthesis for "guided AI" in cybersecurity education, with a focus on beginner-level CTF-based assessments.

Guided AI refers to structured, scaffolded interactions in which students (1) plan before prompting, (2) request hints instead of solutions, and (3) justify or critique AI-generated suggestions. A narrative review of 28 peer-reviewed sources across cybersecurity education, AI in education, and responsible AI use finds that AI support often improves efficiency and first-pass correctness but less often strengthens deeper reasoning, long-term retention, or transfer. Few studies explicitly compare guided versus unguided AI in cybersecurity contexts, and almost none measure trust calibration, over-reliance, or misuse beyond self-report. These gaps motivate a framework that organizes guided AI designs along three dimensions: (a) when help is available (for example, after students articulate an initial claim or plan), (b) how help is constrained (for example, hint-first, error-diagnosis, or explanation-only modes), and (c) what evidence students must provide to demonstrate that they—not the model—own the reasoning.

The lightning talk will highlight concrete patterns, such as (1) requiring students to submit a brief claim-evidence-warrant (CEW) justification before unlocking an AI hint, (2) logging AI interactions alongside CTF submission traces, and (3) designing "AI-visible" versus "AI-resilient" challenges that make reasoning steps observable for both assessment and research. It will also outline reporting standards for future studies, including minimal descriptions of AI configuration, guardrails, data collection, and integrity controls, so that results across courses and institutions become comparable. The session will close with an overview of the future work that will include piloting a Socratic guided-AI application and a practical roadmap based on the research that can be adapted to other CTF assessments in cybersecurity education.

Keywords: Guided AI, cybersecurity education, capture-the-flag, generative AI, large language models, academic integrity, learning analytics.

Engaging the next generation: Cybersecurity outreach for high school students at St. Mary's University

[Lightning Talk]

Ayad Barsoum, St. Mary's University, TX, abarsoum@stmarytx.edu

Extended Abstract

For the 2025 NCAE-C annual requirements report, designated institutions are required to document outreach activities targeting K–12 students. In this lightning talk, we will highlight the ongoing commitment of St. Mary's University in San Antonio, Texas, to STEM engagement through our Tech Camps program, which offers ten technology-focused camps for high school students. These camps reached a diverse population, including students from multiple racial and ethnic backgrounds, students with disabilities, and those from low-income families. Many school districts in San Antonio lack the resources or qualified teachers to teach cybersecurity, underscoring the critical role of St. Mary's University in supporting these students by offering this summer program free of charge.

Of the 10 camps, two were dedicated to cybersecurity education, including one exclusively for female students to encourage greater gender representation in the field. The camps introduced participants to cybersecurity principles, ethics and cybercrime, security threats and attacks, secure internet practices, social engineering, fundamentals of computer networking, and cybersecurity-related career opportunities. A central goal was to increase awareness of post-secondary cybersecurity education and inspire students to consider college programs and careers in cybersecurity.

Program enrollment reached 188 students, exceeding the target of 180, with an 85% overall completion rate above our 80% target. Participant outcomes were systematically assessed using pre- and post-camp surveys, which measured awareness, interest, and intent regarding cybersecurity education and careers. Survey results indicate measurable gains: interest in pursuing a cybersecurity-related college degree rose from approximately 70% to 90%, and awareness of cybersecurity career opportunities increased from 80% to 100%. These findings demonstrate that St. Mary's University's outreach activities effectively enhance student awareness of post-secondary cybersecurity education and career pathways, particularly among underrepresented populations, supporting the objectives of both the NSA NCAE program and broader STEM education initiatives.

Keywords: Cybersecurity education, K–12 outreach, cybersecurity workforce development.

AI-2027 prediction and the future of cybersecurity education

[Lightning Talk]

Marufu Lamidi, Century College, MN, marufu.lamidi@century.edu

Extended Abstract

The AI-2027 project speculates or predicts that artificial intelligence systems may rapidly progress toward human-level or superhuman capabilities within the next decade. Despite uncertainties in the timeline and assumptions, this project raises cybersecurity education concerns due to a curricula gap in preparing cybersecurity students for a future that could be transformed by AI-driven cyber risks. Advances in AI could accelerate vulnerability discovery, automate the development of cyberattacks, enhance large-scale social engineering, and facilitate sophisticated misinformation campaigns. If such capabilities become available to threat actors, cybersecurity professionals will need additional skills to defend these increasingly intelligent, AI-autonomous systems.

This presentation examines how emerging AI development, as highlighted by the speculative AI-2027 project, should inform practical changes to cybersecurity education curricula and teaching approaches to better prepare the future cybersecurity workforce. It identifies emerging competency gaps in AI security, adversarial machine learning, and AI governance. The presentation discusses actionable curriculum adaptations, including integrating AI security modules, developing adversarial AI laboratory exercises, and aligning with workforce competency frameworks. Additionally, pilot implementation strategies are discussed to help institutions incrementally adapt cybersecurity education to an AI-enabled threat environment.

Keywords: Curriculum transformation, future workforce readiness, AI foresight, AI-drive threats.

References:

Kokotajlo, D., Alexander, S., Larsen, T., Lifland, E., & Dean, R. (2025). *AI 2027: Scenario-based forecasting for transformative AI*. AI Futures Project. https://ai-2027.com/

From workforce training to research incubation

[Lightning Talk]

Ulku Clark, University of North Carolina Wilmington, NC, clarku@uncw.edu

Bilge Karabacak, University of North Carolina Wilmington, NC, karabacakb@uncw.edu

Geoff Stoker, University of North Carolina Wilmington, NC, stokerg@uncw.edu

Aysun Karamustafaoglu, University of North Carolina Wilmington, NC, ak2756@uncw.edu

Extended Abstract

Operational Technology (OT) systems are foundational to critical infrastructure sectors such as maritime transportation. However, many academic institutions face challenges in providing sustained hands-on training in OT cybersecurity. High laboratory costs, safety and operational constraints, interdisciplinary skill requirements, and limited student enrollment make it difficult to maintain OT cybersecurity programs. These barriers remain despite the growing demand for professionals who can secure industrial control systems (ICS), programmable logic controllers (PLCs), and cyber-physical maritime systems that connect navigation, vessel monitoring, and operational networks.

This lightning talk presents lessons learned from a research-oriented maritime OT cybersecurity hackathon that was designed to address these challenges. The three-day event was funded through a regional workforce development initiative and brought together students, faculty, industry practitioners, and federal government representatives. The event combined focused foundational instruction with hands-on work using real maritime IT/OT systems, followed by team-based challenges based on realistic OT security scenarios and operational constraints. This hackathon emphasized cyber-physical system behavior and operational constraints present in real maritime OT systems. One challenge was intentionally designed to generate research questions and inform future grant development. Participants analyzed vulnerabilities, cyber-physical dependencies, and operational constraints in the maritime technologies studied during the event. The documented observations later informed future applied research and grant proposal development.

Drawing on participant feedback and instructor observations, the talk highlights the impact of the hackathon's short-format and immersive design. The event helped participants from diverse backgrounds quickly align, better understand cyber-physical constraints unique to maritime systems, and identify early-stage research problems. Beyond skill development, the event also served as a research incubator. It revealed feasibility limits, data gaps, and coordination challenges that are often difficult to identify through traditional coursework or individual research efforts.

The presentation concludes with practical takeaways for CAE-CD institutions interested in using research-oriented hackathons as a scalable complement to formal programs. Lessons learned include resource considerations, the industry partnerships needed to access operational technologies, and how short-duration events can support both hands-on OT cybersecurity learning and early-stage research development in specialized critical infrastructure domains.

Keywords: Maritime cybersecurity, OT security, hackathons, applied research.

Maximizing student potential through student club activities

[Lightning Talk]

Tahir M. Khan, Western Illinois University, IL, tm-khan@wiu.edu

Extended Abstract

Institutions encourage students to participate in various activities, such as conducting peer trainings, organizing Capture the Flag (CTF) events, preparing and presenting materials related to cybersecurity awareness and other topics of interest, and participating in student- or university-organized events. This presentation emphasizes strategies that can be used to maximize student potential by engaging them through club activities. A club advisor can create opportunities and provide a platform for students to lead initiatives, helping them build confidence and leadership skills while completing their degree programs.

As an advisor of a cybersecurity club, I have strategically involved students—particularly executive team members—in both the planning and execution phases of activities. These efforts include, but are not limited to, encouraging students to contact local community organizations to organize awareness trainings; coordinating with secondary education schools to conduct sessions on Artificial Intelligence (AI) and other topics of interest; organizing CTF events for university students; participating in online CTF competitions; assisting students in preparing for professional certifications; and preparing and delivering lessons that help students develop new competencies outside the classroom.

Notable initiatives that were partially or fully coordinated and organized by students, with the advisor support, include: 1) making arrangements with City Hall administrators to organize a training session for executive members, attended by the City Administrator, Community Development Director, Finance and HR Director, and others; 2) coordinating with local school officials to organize a session helping students understand the impact of AI on future careers and businesses; 3) inviting external cybersecurity professionals to share their expertise; and 4) organizing and participating in CTF events. An advisor can implement strategic planning approaches to engage students in activities that build confidence and initiative while improving public speaking, coordination and organizational skills, time management, and professional interaction. In combination, these skills help students maximize their potential and positively impact their future careers.

Keywords: Club activities, club advisor, maximizing student potential, student activities and engagement, lead initiatives.

Rethinking student assessment in the era of artificial intelligence

[Lightning Talk]

Tahir M. Khan, Western Illinois University, IL, tm-khan@wiu.edu

Abdul Salam, State University of New York Polytechnic Institute, NY, salama@sunypoly.edu

Extended Abstract

The use of Artificial Intelligence (AI), particularly generative AI technologies, has increased significantly in recent years. AI tools such as Chat Generative Pre-trained Transformer (ChatGPT), Google Gemini, Microsoft Copilot, and similar applications have been developed to assist with daily tasks, including proofreading content, paraphrasing to improve sentence flow, and performing other functions that enhance efficiency in routine activities.

AI technologies are offered through both paid and free subscription models. Instructors increasingly leverage these tools to create instructional activities and examples that support effective explanation of course content in the classroom. Conversely, students use these tools to prepare for assessments and assist with completing assignments.

Policies governing the use of AI tools vary across institutions. In some cases, the use of AI is left to the discretion of the instructor, who may choose to allow or prohibit its use within a course. Some instructors argue that students may complete assignments using AI without fully understanding the underlying content. To address this concern, instructors have reverted to more traditional assessment methods. For example, reports indicate that faculty members have begun administering oral examinations or paper-based exams without the use of computers to better assess student understanding and assign grades more fairly. Additionally, instructors have implemented project-based learning approaches, requiring student teams to record videos explaining how a project was completed and summarizing the lessons learned.

Although AI writing detection tools are available to help educators evaluate submitted assignments, reports indicate that these tools may not reliably determine whether content was generated by an AI application or written by a human who used the tool only for proofreading or making minor edits to the assigned work. Instructors must therefore strive to use a variety of assessment methods to fairly evaluate student performance. One approach is to use AI tools to generate or attempt the same assessments assigned to students; if the tool can complete the assignment successfully, instructors may need to redesign the assessment to require substantial human reasoning or creativity. Another approach involves brief oral explanations in which students demonstrate their understanding of an assessment; however, this method may disadvantage students facing language barriers. Additional strategies include project-based assessments, oral examinations, and in-class paper-based exams, where feasible.

Keywords: Student assessment, artificial intelligence tools, assessment methods, student performance, rethinking student assessment.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Lessons learned from building an interdisciplinary cybersecurity seminar

[Lightning Talk]

Kevin Floyd, Middle Georgia State University, GA, kevin.floyd@mga.edu

Alan Stines, Middle Georgia State University, GA, alan.stines@mga.edu

Extended Abstract

Cybersecurity challenges increasingly extend beyond technical domains, impacting healthcare, finance, education, public administration, and nearly every discipline that relies on digital infrastructure. In response to this reality, Middle Georgia State University hosted *Cybersecurity Takes a Village: Interdisciplinary Perspectives* in October 2025, a half-day conference that brought together faculty, students, and practitioners from across non-technical and technical disciplines to explore shared cybersecurity challenges and responsibilities.

This lightning talk presents key lessons learned from the design, organization, and execution of an interdisciplinary cybersecurity conference in a higher education setting. Rather than focusing solely on technical content, the event emphasized accessibility, relevance, and cross-disciplinary dialogue, encouraging participants from diverse academic backgrounds to engage meaningfully in cybersecurity topics. The presentation will highlight practical insights related to stakeholder engagement, framing cybersecurity concepts for non-technical audiences, balancing academic rigor with approachability, and fostering productive discussions across disciplinary boundaries.

The session will also discuss challenges encountered during planning and delivery, including aligning expectations across disciplines, avoiding overly technical language, and measuring impact beyond traditional cybersecurity metrics. Attendees will be invited to reflect on how similar interdisciplinary approaches can strengthen cybersecurity education, awareness, and collaboration within their own institutions and CAE-CD communities. By sharing these experiences, this lightning talk aims to spark discussion around scalable, interdisciplinary models for cybersecurity engagement and to encourage the CAE-CD community to view cybersecurity not solely as a technical discipline, but as a shared institutional and societal responsibility.

Keywords: Interdisciplinary cybersecurity, cybersecurity awareness, collaboration, engagement.

Beyond grades: Authentic assessment in cybersecurity education

[Lightning Talk]

Hondo Tamez, Johnson County Community College, KS, atamez@jccc.edu

Extended Abstract

Cybersecurity and information technology education face a persistent challenge: traditional assessments often measure theoretical knowledge without validating the practical skills required to secure systems and respond to real threats. In an environment where breaches can disrupt critical infrastructure and compromise sensitive data, assessment methods must move beyond exams and toward demonstrations of operational competency.

This lightning talk explores practical approaches for integrating authentic assessment into IT and cybersecurity courses by aligning classroom activities with workforce expectations and competencies identified in the NICE Cybersecurity Workforce Framework. Rather than relying primarily on quizzes or written exams, students can demonstrate their understanding through applied security scenarios that mirror tasks performed by cybersecurity professionals.

Examples include system hardening exercises in cloud or virtualized environments where students identify and remediate security misconfigurations, incident response scenarios that require analyzing alerts and recommending response actions, and governance or compliance reviews that ask students to evaluate organizational security practices. These activities shift assessment toward evaluating investigation, problem solving, and communication rather than memorization alone.

The session will share practical strategies for designing assessments that validate real security skills while remaining feasible for instructors to implement in existing IT and cybersecurity courses. By reframing assessment as a tool for skill validation rather than grade assignment, educators can strengthen the connection between academic preparation and workforce readiness.

Keywords: Cybersecurity education, authentic assessment, workforce readiness, NICE Framework, experiential learning.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Teaching cybersecurity policy through cyber policy competitions: A novel approach for technical undergraduate programs

[Lightning Talk]

Stu Steiner, Eastern Washington University, WA, ssteiner@ewu.edu

Chris Cain, Eastern Washington University, WA, ccain7@ewu.edu

James Headley, Eastern Washington University, WA, jheadley@ewu.edu

Antonio M. Espinoza, Eastern Washington University, WA, aespinoza17@ewu.edu

Extended Abstract

Traditional cybersecurity (cyber) curriculum in technical undergraduate programs emphasize both defensive and offensive techniques, threat analysis, and system exploitation; yet often fail to adequately prepare students for the policy dimensions from real-world cyber incidents. While case study analysis in a policy and law course provides foundational knowledge, students rarely transition from understanding precedent to actively shaping policy in complex, evolving scenarios. This talk presents a competition-based pedagogical framework implemented across three national security simulation exercises for undergraduate students.

Uniquely, teams combine both cybersecurity students and political science students, mirroring the interdisciplinary collaboration required in actual government and private sector cyber response. Each competition presents participants with realistic, evolving crisis scenarios involving nation-state-affiliated threat actors deploying attacks against critical systems. Student teams must develop comprehensive response strategies that address both immediate and long-term technical remediation and policy frameworks to prevent recurrence. Their recommendations are delivered in a two-page policy paper, a one-page decision document, and an oral briefing where the students defend their recommendations to a simulated National Security Council panel composed of expert judges from industry, government, and academia.

Our approach systematically develops five underemphasized core competencies critical to cyber professionals: (1) policy analysis capabilities; (2) strategic thinking for developing forward-looking, actionable policies; (3) crisis management skills for responding to dynamic national security scenarios; (4) effective communication of policy and technical analyses; and (5) multidisciplinary integration connecting technical, political, economic, and diplomatic dimensions. The interdisciplinary team structure proves particularly valuable. Cybersecurity students gain fluency in policy frameworks, diplomatic considerations, and stakeholder analysis, while political science students develop technical literacy necessary to craft informed, implementable policies. This interdisciplinary approach creates rich peer learning opportunities and prepares graduates for the collaborative, boundary-spanning roles increasingly demanded in cybersecurity leadership.

Keywords: Policy, competitions, policy simulation, interdisciplinary learning, crisis communication.

Bridging the experience gap: Scaling micro-internships for cybersecurity students through industry partnerships

[Lightning Talk]

Jason Hammon, Western Governors University, UT, jason.hammon@wgu.edu

Extended Abstract

The cybersecurity workforce faces a paradox: over 514,000 job openings exist nationally with only 74% filled, yet entry-level candidates struggle to secure their first role due to experience requirements (CompTIA et al., 2025; Lightcast, 2024). To address this gap, WGU piloted micro-internships through Riipen, a project-based learning platform connecting cybersecurity students with real industry projects. Virtual teams of 5 students tackled cybersecurity challenges over 8 weeks for a total of 40-60 hours. Rather than building industry partnerships directly, WGU leveraged Riipen to source and manage employer connections at scale. This enabled asynchronous participation for a nationally distributed student population and removed typical barriers of geography, scheduling, and limited professional networks.

Students applied through a self-selected marketplace after completing a defined sequence of foundational courses, ensuring a consistent competency floor while allowing teams to leverage complementary strengths. Projects centered on applied cybersecurity work including risk assessments aligned to NIST/ISO frameworks, compliance readiness, secure cloud configuration, email security diagnostics, and platform security QA — each producing professional deliverables for real organizations across healthcare, fintech, nonprofit, and SaaS sectors.

Results from 65 enrolled students were encouraging: 74% completed the program, and over 80% reported improvements in collaboration, communication, and problem-solving. Critically, 84% applied coursework skills to real projects, 74% felt better prepared for employment, and 68% saw themselves as more marketable. All participating employers said they would hire these students if positions opened, rating teams 4.97/5 on average. Students received digital badges and could post employer feedback to LinkedIn to signal to future hiring managers their "real world" experiences. This model shows how institutions can scale meaningful industry experiences through partnerships. Attendees will leave with actionable implementation guidance, including how to scope projects for student success, structure employer onboarding, leverage a marketplace model to increase student motivation, and identify the right prerequisite baseline for participation.

Keywords: Cybersecurity careers, experiential learning, internships, industry partnerships.

References:

CompTIA, Lightcast, & National Initiative for Cybersecurity Education. (2025). *Cybersecurity supply and demand heat map* [Interactive data visualization]. CyberSeek. https://www.cyberseek.org/heatmap.html

Lightcast. (2024, October). *The cybersecurity gap: Quarterly white house report Q3 2024*. https://lightcast.io/resources/research/quarterly-cybersecurity-talent-report-oct-24

AI fluency as a cross-disciplinary cybersecurity skill

[Lightning Talk]

Elliott S. Lynn, American Public University System, WV, Elliott.lynn@mycampus.apus.edu

Extended Abstract

Cybersecurity professionals increasingly operate within environments that span technical systems, organizational processes, legal frameworks, and human behavior. As artificial intelligence tools become embedded across these domains, the ability to work effectively with AI is no longer confined to purely technical roles. Instead, AI fluency is emerging as a cross-disciplinary skill that supports collaboration, decision-making, and risk management across cybersecurity and related fields.

This lightning talk examines AI fluency as a shared competency that connects cybersecurity with adjacent disciplines such as data analytics, governance and compliance, systems engineering, and policy. Rather than focusing on specific tools or technologies, the session emphasizes the underlying skills required to engage with AI-enabled systems responsibly and effectively. These skills include articulating domain-specific questions, evaluating AI-generated outputs, recognizing limitations and bias, and applying professional judgment within ethical and regulatory constraints.

The talk highlights how cybersecurity education can intentionally cultivate AI fluency as a transferable skill that enhances interdisciplinary collaboration. By integrating AI-enabled activities into coursework that spans technical, managerial, and policy-oriented perspectives, programs can better prepare students to function in environments where cybersecurity decisions are informed by multiple stakeholders and data sources. This approach reinforces the role of cybersecurity professionals as integrators of technology, risk, and organizational context rather than isolated technical specialists. By framing AI fluency as a cross-disciplinary capability, this session contributes to CAE-CD discussions on workforce readiness and educational alignment. Attendees will gain a concise perspective on how emphasizing shared AI skills can strengthen cybersecurity education and better reflect the realities of modern cyber defense practice.

Keywords: Artificial intelligence fluency, interdisciplinary cybersecurity, workforce readiness, cybersecurity education, professional skills.

Beyond the textbook: Growing defenders in a student security operations center

[Lightning Talk]

Michael Ramage, Murray State University, KY, mramage@murraystate.edu

Randall Joyce, Murray State University, KY, rjoyce@murraystate.edu

Extended Abstract

With the ever-increasing need for cybersecurity professionals with practical experience, students must acquire hands-on skills alongside their theoretical learning. Through a program funded by the U.S. Department of Labor, our university established a SOC (security operations center) that provides students with real-world experience in this environment. In this environment, both in-person and online students at the undergraduate and graduate levels are working with an Elastic Stack SIEM (security information and event management) system to provide additional review of traffic logs for the campus information security team. Students can participate in the SOC through Threat Hunting and SOC Management, an Independent Study, and an Internship Experience course. The Threat Hunting and SOC Management courses and the Independent Study are offered in our undergraduate program. At the same time, the Internship Experience was developed for the U.S. Department of Labor grant focused on our cyber analyst certificate program. Through these avenues for students to get involved in the student SOC, they learn and practice the skills of a cybersecurity analyst, including threat intelligence, SIEM operations, network protocol review, and incident response, all the while using real-world data to learn from and using our expertise to give back. At the same time, they are improving their soft skills (communication, problem-solving, critical thinking) through daily operations in the SOC. Providing students with this opportunity during their educational journey helps prepare them for the workforce and supports the university. The student-centered SOC continues to grow in two key ways. First, a physical SOC is being built to provide a real in-person SOC experience. Second, communities and organizations from across the state are requesting our help. In the future, our students will be able to support communities and organizations statewide that cannot afford cybersecurity monitoring services.

Keywords: Security operations center, security information and event management, career development, cybersecurity analyst, threat intelligence.

Using locally installed LLMs to support ethical hacking and cybersecurity exploration in a controlled environment

[Lightning Talk]

Kevin Lann-Teubner, Butler Community College, KS, klannteubner@butlercc.edu

Brian Dye, Butler Community College, KS, bdye@butlercc.edu

Extended Abstract

As AI systems become more influential in cybersecurity, educators and researchers must navigate the challenge of using these tools within environments that require flexibility, transparency, and realistic scenarios. Cloud based LLM are limited by strict safety standards, inconsistent behavior, and dependency on external services. This presentation explores how locally installed large language models provide controlled, institutionally managed alternatives for teaching and learning ethical cybersecurity topics in community college courses and a student cybersecurity club.

By running models locally, cybersecurity programs gain the ability to craft customized lab exercises, simulate adversarial conditions, and generate instructional material without exposing students or networks to uncontrolled or unethical content. Lessons learned from implementation will be shared. Student focused lessons include the importance of teaching effective AI prompting, setting realistic project scope, and focusing on proof-of-concept experimentation. Administrative considerations include planning for hardware requirements, addressing institutional security concerns, and adapting to the rapid pace of change in the AI ecosystem.

Keywords: Locally deployed LLMs, AI-assisted security training, controlled AI environments, responsible AI use.

Building a sustainable cybersecurity teacher pipeline

[Lightning Talk]

Jenny Ju, City University of Seattle, WA, jujenny@cityu.edu

Morgan Zantua, City University of Seattle, WA, zantuamorgan@cityu.edu

Extended Abstract

The national cybersecurity workforce shortage underscores the need for earlier and more inclusive cybersecurity education pathways. While many initiatives focus on high school and postsecondary programs, middle school is a critical inflection point for shaping student interest, self-efficacy, and career awareness. Learning About Techniques and Tools for E-security (LATTE) is a proposed initiative designed to address this gap by preparing middle school teachers to integrate cybersecurity concepts across academic subjects and extracurricular activities.

Building on the proven Cybersecurity High School Innovations (CHI) model, LATTE extends the cybersecurity education pipeline earlier by leveraging CHI's alumni network and curriculum expertise. LATTE delivers a capacity-building model that combines online professional development with an in-person Summer Summit and awards Continuing Education Units (CEUs).

LATTE employs a socio-technical approach to cybersecurity education that provides middle school teachers with curriculum design skills that integrate active learning and continuous assessment. Teachers engage with modular, open-access curriculum resources from cyber.org, teachcyber.org, CLARK, RING, and PICO, and learn strategies to embed cybersecurity topics into disciplines such as mathematics, science, language arts, and social studies. Experiential learning is reinforced through after-school clubs and cyber defense competitions, supported by coaching from CHI-trained educators, industry professionals, and National Guard partners.

By equipping middle school teachers with content knowledge and pedagogical strategies, LATTE establishes a repeatable and sustainable model for broadening participation and strengthening the K–12 cybersecurity education pipeline.

Keywords: Cybersecurity education, capacity building, teacher professional development, socio-technical learning.

Volunteering as applied cybersecurity education

[Lightning Talk]

Christopher Kadlec, Georgia Southern University, GA, ckadlec@georgiasouthern.edu

Elizabeth Rasnick, University of West Florida, FL, erasnick@uwf.edu

Extended Abstract

The cybersecurity industry continues to grapple with a paradoxical "experience gap": entry-level candidates possess the theoretical knowledge and certifications required for employment but lack the exposure to live, "high-entropy" environments necessary for professional efficacy. While traditional academic labs and Capture the Flag (CTF) competitions provide valuable technical scaffolding, they are often sanitized, predictable, and devoid of human-centric friction.

This presentation proposes Volunteering and Community Service-Learning as a critical "Third Pillar" of applied education, sitting between formal classroom instruction and professional internships. By engaging with under-resourced non-profits, small businesses, and local government entities, students are forced to navigate the messy reality of legacy systems, limited budgets, and the "human firewall."

The Pedagogical Framework

We argue that volunteering serves as a unique pedagogical laboratory where the NICE Workforce Framework for Cybersecurity (SP 800-181) comes to life. Unlike controlled simulations, community service requires students to perform:

- **Vulnerability Management** on systems they did not build and cannot easily replace.
- **Governance and Risk Management (GRC)** for organizations with zero existing security policy.
- **Security Awareness Training** for non-technical users, requiring a mastery of soft skills and translation of technical jargon.

Institutional & Student Benefits

For CAE-designated institutions, this fulfills the "Community Outreach" and "Professional Development" requirements of the CAE designation while strengthening the local cyber ecosystem.

This talk challenges the CAE community to move beyond the sandbox. By integrating volunteering into the curriculum—whether through service-learning credits, capstone projects, or student-led clinics—we can protect our most vulnerable community assets while simultaneously forging the next generation of highly-adaptive, socially-conscious cybersecurity professionals.

Keywords: Volunteerism, applied education, internships, cybersecurity, pedagogy.

Using threat modeling to teach introductory cybersecurity

[Lightning Talk]

Frederick Scholl, Quinnipiac University, CT, frederick.scholl@quinnipiac.edu

Morrow Long, Quinnipiac University, CT, harry.long@quinnipiac.edu

Extended Abstract

We have developed and used methodologies and course materials to teach threat modeling in our introductory cybersecurity course. This approach offers a holistic view of cybersecurity while addressing the requirements of multiple CAE-CD KU's. In this presentation we will present our approaches and share our course materials.

"Introduction to Cybersecurity" is the first course in our CAE-CD validated MS Cybersecurity degree program. The challenge in this course is to provide a big picture of security processes and technology, while meeting detailed KU requirements. "Threat modeling" was originally developed to improve software security; it is still widely used in this application. For our class we apply it to problems of system risk assessment and risk management. Source materials are mini-case studies generated partly by AI (Artificial Intelligence) LLM's (Large Language Model) with input from real breaches. The classroom exercise includes enumeration of system assets, and identification of threats, vulnerabilities and risks.

We use two threat modeling methods in the class. The first is based on the Elevation of Privilege (EOP) game (Shostack, 2014). Students are divided into groups of six and given a system description, playing cards and game instructions. We created new playing cards including threats and mitigations. Each round lasts approximately one hour at which time students report to the class on their findings. The benefit of this method is that students learn a hands-on threat modeling approach and start to describe real system risks.

Our second approach makes use of the University Edition of a commercial SaaS tool from Threat Modeler (https://www.threatmodeler.com). Students are again given systems descriptions from which they build suitable models. The application outputs both threats and security requirements. Students address infrastructure, security threats, and risks as well as security mitigations. We will demonstrate input artifacts, the application user interface and output reports.

These two exercises address several KU's and KU topics. These include: SRA, CSP, CSF and ISC. Mappings will be provided. These topics are then reinforced in subsequent classes.

In summary, our approach to System Threat Modeling provides hands-on classroom experiences, at no cost or low cost, while addressing business needs for cybersecurity and the CAE-CD KU's. These approaches provide a "big picture" perspective of cybersecurity while utilizing techniques that can be employed by students in professional settings.

Keywords: Threat modeling, risk management, risk assessment, introductory cybersecurity, KU.

References:

Shostack, A. (2014). *Threat modeling: designing for security.* Wiley.

Toward experiential training program for AI security and privacy practitioners

[Lightning Talk]

Mujtaba Nazari, Loyola University Chicago, IL, mnazari@luc.edu

Mohammed Abuhamad, Loyola University Chicago, IL, mabuhamad@luc.edu

Loretta Stalans, Loyola University Chicago, IL, lstalan@luc.edu

Eric Chan-Tin, Loyola University Chicago, IL, dchantin@luc.edu

Extended Abstract

The rapid adoption of Artificial Intelligence across industries has outpaced security and privacy training for AI practitioners. This talk presents methods, modules, and findings from an experiential training program designed to address security and privacy challenges in AI systems development and deployment. We conducted two program iterations: a comprehensive 12-workshop series (May-October 2024) and a condensed 6-workshop format (January-February 2025). The program combined expert-led panel sessions with hands-on laboratory activities, engaging 78 participants from diverse professional backgrounds. Evaluation through pre- and post-evaluation surveys and qualitative observations revealed improvements in cybersecurity knowledge and AI security awareness. Participants demonstrated enhanced ability to identify vulnerabilities, implement security measures, and develop organizational policies for AI-related risk mitigation. The condensed format showed comparable learning outcomes with improved completion rates. This effort highlights the increased need to establish cybersecurity and privacy training for AI professionals to develop secure and trustworthy AI systems.

This talk was previously presented at The Colloquium for Information Systems Security Education (CISSE) 2025 and a full paper has been accepted for publication at the Journal of The Colloquium for Information Systems Security Education (CISSE). This talk is appropriate for the CAE-CD Symposium as it covers both cybersecurity and AI.

Keywords: AI security, cybersecurity training, privacy-preserving AI, experiential learning, professional development, adversarial machine learning.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

The Nevada Cyber Range (NCR): A scalable platform for cybersecurity education and operations

[Lightning Talk]

Jake Herweg, University of Nevada, Reno, NV, jherweg@unr.edu

Shamik Sengupta, University of Nevada, Reno, NV, ssengupta@unr.edu

Extended Abstract

Delivering effective instruction in cybersecurity and computer science requires access to computing environments that realistically reflect operational systems while preserving the security of institutional networks and sensitive user data. Traditional physical laboratory infrastructures composed of networked machines present significant limitations, including high capital and maintenance costs, limited scalability, and constraints on administrative access. These factors hinder the ability to provide flexible, hands-on learning environments that support authentic, real-world problem solving. Although virtualized environments offer a viable alternative, reliance on third-party cloud platforms can be cost-prohibitive, while building and maintaining private virtual infrastructures demands substantial technical expertise and ongoing administrative overhead.

To address these challenges, we have developed the Nevada Cyber Range (NCR), a scalable, virtualization-based instructional platform built on Proxmox. NCR provides students with isolated, live virtual machines (VMs) running multiple operating systems, over which they are granted full administrative control. This design enables learners to perform advanced configuration and networking tasks, including firewall rule development, routing, traffic analysis, scanning, and logging, within a secure and controlled environment. Each student is provisioned with dedicated VMs that can be flexibly arranged into custom network topologies tailored to specific instructional objectives. To further enhance accessibility and scalability, the platform utilizes a browser-based proxy, removing the need for virtual private network (VPN) connections, while Ansible automation streamlines the deployment of complex, isolated network topologies for large classes. For Linux-based systems, individual VMs may host internal networks of containerized instances, allowing students to interact with and analyze live, multi-node networks rather than abstract simulations.

Recently, NCR expanded its scope beyond education to include a functional Security Operations Center (SOC) through a partnership with IBM and CarbonHelix. This collaboration has enabled the deployment of IBM QRadar, an enterprise-grade security information and event management (SIEM) platform, to monitor traffic and identify anomalies across both physical labs and virtual NCR environments. This integration not only strengthens the university's security posture but also provides students with hands-on experience in SOC analyst workflows using industry-standard tools. Future initiatives aim to extend this monitoring support to under-resourced local governments in Nevada, fostering a statewide cybersecurity ecosystem driven by student expertise.

Keywords: Cybersecurity education, virtualization, security operations center, SIEM, hands-on learning.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

CPPJ - Cybersecurity Pedagogy and Practice Journal: Origins, evolution, purpose

[Lightning Talk]

Anthony Serapiglia, Saint Vincent College, PA, anthony.serapiglia@stvincent.edu

Extended Abstract

Cybersecurity education has matured rapidly over the past decade, yet much of the scholarly conversation in the field continues to emphasize technical novelty over instructional design, experiential learning, and evidence-based teaching practice. Faculty working within Centers of Academic Excellence (CAE) programs routinely develop innovative laboratories, curricula, and assessment models, but these contributions often lack a dedicated scholarly venue aligned with the realities of cybersecurity education. The Cybersecurity Pedagogy and Practice Journal (CPPJ) was founded to address this gap.

This lightning talk presents the origins, mission, and scope of CPPJ as an open-access, peer-reviewed journal dedicated specifically to the scholarship of teaching and practice in cybersecurity. The journal was established to provide a rigorous yet accessible publication outlet for faculty and practitioners who design, implement, and evaluate cybersecurity learning experiences across higher education, training programs, and workforce development initiatives. Unlike traditional security journals, CPPJ prioritizes pedagogical relevance, reproducibility of instructional approaches, and the translation of classroom and lab innovations into transferable practice.

The presentation briefly traces the motivations behind CPPJ's founding, including recurring challenges faced by CAE-affiliated educators: limited venues for publishing applied teaching research, misalignment between educational contributions and traditional technical review criteria, and the need for community-centered scholarship that values instructional impact. CPPJ's editorial model, review process, and topical focus will be discussed to illustrate how the journal supports cybersecurity education while maintaining academic rigor.

For the CAE-CD community, CPPJ represents a mechanism to capture and disseminate effective practices, lessons learned, and pedagogical innovations that strengthen cybersecurity programs beyond lesson plan repositories. This lightning talk aims to raise awareness of CPPJ, clarify its role within the broader cybersecurity scholarship ecosystem, and encourage CAE educators to contribute their teaching-focused work to a venue designed expressly for their needs.

Keywords: Cybersecurity education, pedagogy, publishing, scholarship of teaching, CAE community.

References:

Information Systems and Computing Academic Professionals. (2022). *Cybersecurity Pedagogy and Practice Journal.* https://cppj.info/about.html

Building cybersecurity talent through apprenticeship: A success story from the community college of Baltimore county and state employer

[Lightning Talk]

Noell Damron, Community College of Baltimore County, MD, cdamron@ccbcmd.edu

Vini Nithianandam, Community College of Baltimore County, MD, vnithiana@ccbcmd.edu

Sabum Anyangwe, Community College of Baltimore County, MD**,** sanyangwe@ccbcmd.edu

Extended Abstract

Maryland faces a growing demand for skilled professionals in cybersecurity and other high-demand fields. The Community College of Baltimore County (CCBC) Apprenticeship Center helps meet this need through registered apprenticeship programs that combine paid, on-the-job training with classroom instruction. The Community College of Baltimore County (CCBC) Apprenticeship Center in partnership with the State Employer, CCBC is addressing critical workforce gaps by creating sustainable, debt-free pathways into careers such as cybersecurity. Our Cybersecurity Support Technician Apprenticeship exemplifies how paid, hands-on training combined with classroom instruction creates a sustainable talent pipeline while offering apprentices a debt-free pathway to career success.

CCBC's sponsored programs are designed with input by and for multiple employers The Cybersecurity Support Technician Apprenticeship allows participants to earn wages from day one while gaining hands-on experience, theoretical knowledge, and mentorship. Apprentices complete specialized training in network security fundamentals, manage firewalls, incident response and risk mitigation, and system hardening and vulnerability management. Upon completion, they earn CCBC credentials and a Maryland Department of Labor journeyperson certificate, opening doors to advanced roles, entrepreneurship, and global career opportunities.

The program transforms lives and careers. Apprentices describe the program as "The apprenticeship has given me the opportunity to exercise so many things I've learned in cybersecurity. It allows a beautiful blend of creativity and technical skill." "This experience has strengthened my skills and professional growth, and I hope it encourages others to explore a similar pathway." "Success doesn't fall into your lap—you have to reach out, grab it, and secure it just like you would protect a network."

CCBC's apprenticeship programs demonstrate that integrating academic learning with hands-on experience is a powerful solution to workforce shortages. By preparing apprentices for diverse opportunities locally and globally, these programs not only meet employer needs but also empower individuals to thrive in high-demand careers.

Keywords: Cybersecurity, workforce, certification, apprenticeship, education, partnership.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

From 'What can I do in cyber?' to 'Where do I go from here?': Igniting interest with try cyber micro-challenges

[Lightning Talk]

Paige Flores, California State University San Bernardino, CA, paige.zaleppa@csusb.edu

Extended Abstract

Cybersecurity often feels inaccessible to those outside of technical fields—perceived as requiring years of specialized knowledge. Yet the demand for cyber professionals continues to outpace supply, making community outreach critical for the development of a cyber workforce pipeline. This presentation explores how Try Cyber (https://trycyber.com), a no-cost web-based application, can transform community engagement by making cybersecurity tangible, approachable, and exciting for audiences who might never have considered a career in the field.

Try Cyber provides prospective cyber professionals with access to fifteen-minute micro-challenges that simulate day-one internship experiences across nineteen work roles within the NICE Framework and can be easily cross walked to the Department of War Cyber Workforce Framework (DCWF). From Incident Response and Vulnerability Analysis to Digital Forensics and Network Operations, users gain hands-on exposure to real tasks without requiring extensive and often costly technical prerequisites or complex setup. This accessibility makes Try Cyber an ideal tool for schools, career fairs, summer camps, and community events where time is limited and audience technical skills are diverse.

Drawing from experience deploying Try Cyber at two CAE Institutions — engaging students from computer science, business, information technology, health sciences, and mathematics, as well as general community members, — this session will demonstrate practical, replicable strategies for reaching audiences who are curious about the field but unsure where to start. Lessons learned and implementation challenges from both institutions will be shared, giving attendees a realistic picture of what to expect and how to adapt when bringing Try Cyber into their own contexts. Attendees will leave able to identify event formats — such as career fairs, open houses, and summer camps — where Try Cyber's digital mentor-guided challenges can be integrated to spark interest, facilitate work role discovery, and create memorable first experiences with cybersecurity concepts, all without requiring specialized equipment or extensive facilitator training.

Keywords: Community outreach, hands-on learning, career pathways, DCWF, cybersecurity workforce development.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Certified skills list: Translating academic performance into verifiable skills

[Lightning Talk]

Ram Dantu, University of North Texas, TX, ram.dantu@unt.edu

Cihan Tunc, University of North Texas, TX, cihan.tunc@unt.edu

Astro Pryor, University of North Texas, TX, stellapryor@my.unt.edu

Fiho Lee, University of North Texas, TX, fiholee@my.unt.edu

Vinh Quach, University of North Texas, TX, vinh.quach@unt.edu

Extended Abstract

Employers and human resource departments increasingly report that resumes and CVs have become highly standardized in structure and content, making it difficult to distinguish between candidates. Also, traditional academic artifacts, such as transcripts and GPAs, provide limited insight into a student's specific technical abilities. The Certified Skills List (CSL) addresses this gap by translating academic performance into a structured, skills-based credential derived from verified coursework. The CSL aggregates grades from assignments, quizzes, and examinations and maps them to the National Security Agency (NSA)'s Knowledge Units (KUs), producing a standardized representation of validated academic skills – rather than replacing traditional grades, the CSL complements them by highlighting competencies developed through coursework. Unlike self-reported skills or external certification exams, the CSL is generated directly from instructor-evaluated work completed at an accredited institution.

The CSL framework involves three stakeholders: instructors verify assessment data and authorize CSL generation, students select which validated skills are shared, and employers use the CSL as an additional evaluation artifact alongside traditional application materials such as resumes or transcripts. Skills included in the CSL are derived from course learning objectives and assessment outcomes and are cross-referenced with requirements found in job descriptions. For example, coursework in computer networking may validate skills such as network attacks and exploitation, network hardening and network security monitoring. Students retain control over which validated skills appear in their final CSL to present to employers. Skills are identified by a backend large language model (LLM). The fine-tuned model semantically analyzes assessments by converting their text into dense vector representations, allowing it to capture meaning beyond simple keywords. It also filters vague or irrelevant language and, through two learnable gating mechanisms, selects the most appropriate skills based on similarity scores.

To evaluate the practical value of CSL approach, we conducted a survey: 85.72% of attendees found the CSL highly descriptive of candidate abilities, and 71.43% reported that it would be useful in hiring decisions. Initial implementation highlights several lessons, including the importance of clearly defining course-level skills, maintaining instructor oversight of verified competencies, and ensuring student control over shared information.

Keywords: Workforce readiness, knowledge units, skills-based credentialing, academic skill verification, certified skills list.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Cheating or learning? Understanding AI misconceptions in the community college classroom

[Lightning Talk]

Folashade Adeleke, Prince George's Community College, MD, adelekfo@pgcc.edu

Extended Abstract

Rapid advances in generative artificial intelligence have reshaped how community college students learn, yet many instructors continue to view AI primarily as a threat to academic integrity rather than a potential learning partner. Faculty concerns about AI including fears of diminished rigor, unoriginal work, and increased plagiarism are understandable in an environment where tools evolve faster than institutional policy. Yet prohibiting AI rarely prevents its use; instead, it hides it. This talk surfaces three guiding questions that can help instructors reflect on their assumptions and engage in deeper dialogue: **(1) Is the use of AI allowed in your classes? Why or why not? (2) How are your students using AI? (3) How often do use AI? (Professionally or Personally)** These questions encourage transparency and highlight the inconsistency that often emerges when faculty benefit from AI for productivity, content generation, or lesson planning while simultaneously preventing students from using the same tools for learning.

By reframing AI as a literacy rather than a shortcut, instructors can shift their focus toward the skills that truly matter in cybersecurity and related technical fields such as; problem framing, analysis of AI-generated outputs, verification of accuracy, and ethical reasoning. Students who learn to critique and refine AI responses are not avoiding thinking; rather, they are building competencies aligned with real-world cyber defense environments in which AI-assisted tools are becoming standard.

To help instructors see how AI can be integrated without compromising academic integrity, the presentation offers concrete examples of re-designed cybersecurity lab assignments that reduce opportunities for cheating while promoting authentic learning. We can design assignments with academic rigor while acknowledging AI's presence in students' academic lives. These include offering clear AI-use statements, designing assignments that require students to document their process, and incorporating short discussions on evaluating the credibility and limitations of AI outputs. Such approaches help shift faculty from policing to guiding student use, reducing dependence on detection tools and strengthening students' ability to use AI responsibly and critically.

This talk concludes by emphasizing that the most productive question is no longer whether students are using AI—they are—but how instructors can shape that use to promote deeper understanding. When misconceptions are replaced with informed, reflective frameworks, faculty can uphold academic integrity while supporting student success, equity, and workforce readiness in an AI-rich cybersecurity ecosystem.

Keywords: Academic integrity, artificial intelligence (AI), student learning, AI literacy, AI teaching strategies.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Beyond the classroom: Integrating a 24/7 immersive "living & learning" cyber ecosystem

[Lightning Talk]

David Richards, Grand Canyon University, AZ, david.richards1@gcu.edu

Extended Abstract

This session examines Grand Canyon University's "Cybersecurity Living & Learning" initiative, an innovative program aimed at advancing cybersecurity education through the provision of an immersive, continuous, and gamified network environment. The initiative is distinguished by its 24/7 access in residential housing, and the integration of experiential learning methodologies with simulated threat landscapes and cross-disciplinary collaboration, thereby cultivating an ecosystem for sustained skill development. A notable feature of the program is its tiered network design, which permits participants to engage in activities such as network scanning, penetration testing, and other offensive security operations that would typically be prohibited in live production settings.

The platform categorizes network challenges into three levels of complexity: "Easy" networks supply hackable web applications and virtual machines; "Medium" networks emulate corporate infrastructures with relaxed security postures and delayed patching cycles; and "Hard" networks replicate enterprise environments characterized by strong security controls and rigorous patching regimens. Points are allocated in accordance with the difficulty of each breach—introductory exploits yield basic points, while advanced scenarios award greater point values. Learners are required to identify and submit flags as evidence of successful breaches, encouraging a systematic and analytical approach.

All network environments are instrumented with industry-standard security monitoring tools, providing students the opportunity to develop competencies in security operations and incident monitoring. Through direct interaction with these environments, students are able to reinforce theoretical knowledge with practical application, thus narrowing the gap between academic instruction and real-world cybersecurity demands. Attendees will gain insights into the implementation and pedagogical advantages of continuous, experiential learning environments, and recognize the effectiveness of the "living and learning" model in preparing students for the complexities of contemporary cybersecurity practice outside conventional classroom contexts.

Keywords: Student activities, ethical sandbox, competitions, workforce readiness, cybersecurity.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Skills development: Matching certificates to work roles

[Lightning Talk]

Alexis Blackwell, University of North Texas, USA, alexisblackwell@my.unt.edu

Ram Dantu, University of North Texas, USA, ram.dantu@unt.edu

Vinh Quach, University of North Texas, USA, vinh.quach@unt.edu

Extended Abstract

The current status of hiring focuses on being able to accurately and quickly match a candidate with a set of skills to a set of qualifications listed in a job posting. Knowledge, Skills, Abilities, and Tasks (KSATs) are a government defined set of qualifications that cover topics in the Cybersecurity domain. Certain federal jobs are linked to Work Roles, which are directly assigned a list of KSATs that should be achieved. To better meet these required KSATs, job seekers and employers often look towards credentials to provide insight in an individual's set of skills. These credentials are programs that aim to provide training and accreditation for an individual to be proficient in a set of skills and knowledge in a domain.

However, not every credential is mapped to a DCWF Work Role and its associated KSATs, which results in having to manually identify and estimate which credential would provide the correct training to succeed in a Work Role. For hiring managers, the credential summary may not provide an easy-to-understand synopsis of what skills an individual may be proficient in and how those skills match to the job description. For an individual candidate, they may struggle with identifying which certificate would provide the necessary training that they may require to meet a job's requirements.

A methodology has been developed to use a Large Language Model (LLM) to perform the task of automatically matching certificate information to the Work Roles and KSATs. In this process the certificate syllabus and exam text is parsed down into KSATs by the LLM before then matching the certificate KSATs to the Work Roles and the Work Role's KSATs. The end result is a set of Work Roles and a list of KSATs that match to the certificate. Learning taxonomies, such as Bloom's Taxonomy, are also provided to the LLM as a reference for determining the learning level of the content. These learning levels are used to evaluate the KSATs that are considered Basic, Intermediate, or Advanced. In this way, the proficiency level covered by a certificate can also be displayed to the hiring manager and individual. During experimentation, 26 certificates were randomly selected and evaluated for how accurately the LLM could determine if a certificate matched to a Work Role. A 62% accuracy for this methodology was obtained when comparing the LLM's matches to human labeled matches. In conclusion, using this system an individual can select a certificate and then display the matching Work Roles, best matching KSATs, and the percentage of the certificate is considered Basic, Intermediate, and Advanced. Future work will be focused on improving the match percentage and evaluating more certificates.

Keywords: Information matching, work roles, KSATs, large language models.

Scaling student engagement in cybersecurity clubs and competitions

[Lightning Talk]

Leslie Corbo, Utica University, New York, lecorbo@utica.edu

Maxim Gorbachevsky, Utica University, New York, magorbac@utica.edu

Extended Abstract

In 2025, entry-level cybersecurity candidates face a "junior paradox": most postings labeled as entry level now expect substantial, demonstrable hands-on experience (Papadopoulou, 2026). Student cybersecurity clubs and defense competitions have shifted from optional extracurriculars to core work-based learning experiences that often provide the most practical bridge to professional readiness. At the same time, rising technical complexity and uneven support structures in academic programs contribute to participation and persistence gaps, especially for beginners and nontraditional learners.

This lightning talk examines key barriers to student involvement, including tool-centric hiring filters, high perceived entry thresholds, and the lack of structured competency pathways. The NICE Framework Strategic Plan (2021–2025) calls for a shift from primarily theoretical preparation to performance-based assessments, underscoring the need to foreground observable, job-relevant skills (National Institute of Standards and Technology, 2025). The talk also explores how traditional competition formats can inadvertently alienate nontraditional learners and early-career students.

Drawing on current observations within a Cyber Defense University program, the presentation highlights the promise of cybersecurity clinic models—experiential environments that mirror "teaching hospitals" to provide authentic, high-impact learning (Bertone et al., 2025). The session concludes with actionable strategies for 2026, including tiered skill pathways and curriculum-aligned projects that normalize incremental skill development and reinforce the professional value of participation.

Keywords: Cybersecurity workforce readiness, experiential learning, student engagement, cyber defense competitions, entry-level skills gap.

References:

Bertone, B., Wagner, P., & Pauli, J. (2025, September 8). Experiential learning: Innovative approaches to post-secondary cybersecurity education. *Journal of Cybersecurity Education, Research, & Practice, 2025*, 1 (15). https://doi.org/10.62915/2472-2707.1254

National Institute of Standards and Technology. (2025, March 14). *2021-2025 NICE strategic plan.* https://www.nist.gov/itl/applied-cybersecurity/nice/about/strategic-plan

Papadopoulou, E. (2026, January 13). *Cybersecurity readiness report 2026: Same job, new skills. Cyberbit.* https://www.cyberbit.com/cybersecurity-training/cybersecurity-skills-report-2026/

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Here are the missing masses: Centering the home in cyber-education and policy

[Lightning Talk]

Toluwalogo B. Odumosu, Morgan State University, MD,
toluwalogo.odumosu@morgan.edu

Extended Abstract

As recent cybersecurity exploits have shown, home networks are increasingly becoming an attack vector for malicious actors to infiltrate corporate networks. Highlighting the need for reimagining the place and importance of home cybersecurity in cyber-education and cyber-policy. Two recent developments in history have driven significant changes in home networks and heightened the security risks they pose. One is the explosion of IoT devices, driven by the low cost of provisioning everyday devices with 2.4GHz radios, and the other is the rise in work-from-home policies, accelerated during the pandemic, which precipitated a cultural shift to work increasingly undertaken at home. This talk will present an argument for why the CAE community needs to take a critical look at vulnerabilities in Home cybersecurity and will evaluate the risks and challenges of addressing this important domain.

As the 2024 IMF Global Financial Stability Report shows, cyber incidents have been steadily rising, with a significant increase since 2011. The critical importance of home cybersecurity is reflected in NIST research on home routers, which is being driven by recent executive orders. However, there is a critical asymmetry between the vulnerabilities that an inattention to these networks poses and the attendant focus on enterprise infrastructure as represented in the CAE curriculum. Just as the sheer volume of dark matter relative to baryonic matter means that cosmologists can study only portions of the Universe, the Cybersecurity community's inattention to securing home networks may be ignoring the missing masses. In this talk, I will present a compelling case for moving in a new direction concerning cybersecurity in the home and how the community can provide much-needed leadership in this space.

Keywords: Cyber-policy, cyber-education, IoT security, home cybersecurity.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Integrating security & mental health intersection topics in cybersecurity education: A preliminary study

[Lightning Talk]

Ankur Chatterjee, Northern Kentucky University, KY, chattopada1@nku.edu

Michael Otu, Northern Kentucky University, KY, otum1@mymail.nku.edu

Extended Abstract

Cybersecurity professionals face intense mental demands, cognitive stress, high stakes responsibilities, emotional exhaustion, and limited support from organizations. While security breaches and technical defenses dominate public discourse, the mental health of cybersecurity workforce is not given its due importance. A review of recent literature on security and mental health intersections shows the emergence of new subject matters, such as psyber-security, psyber-resilience, psychiatric engineering, cognitive hacking, and indicates a need to study the psychological & emotional impacts (including victim trauma) resulting from cyber incidents.

However, the current cybersecurity curricula, whether it be college-level training or industry-oriented professional certification, lack coverage of topics in the intersection area of security and mental health. We have conducted a survey with people from academia and industry as part of a preliminary study to determine the need plus scope for inclusion of these security and mental health intersection topics into cybersecurity education. This talk will briefly discuss this work-in-progress study and shall also share some initial results plus findings from this study.

Keywords: Cybersecurity education, mental health, review, curricula, survey, study, psyber-security, psyber-resilience.

CD Track: Refereed Extended Abstract Proceedings for Posters

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Important factors in obtaining a cybersecurity job in the United States

[Poster]

Laila Hamdan, Loyola University Chicago, IL, lhamdan@luc.edu

Aslihan Altindal, Loyola University Chicago, IL, aaltindal@luc.edu

Eric Chan-Tin, Loyola University Chicago, IL, dchantin@luc.edu

Extended Abstract

Cybersecurity has become one of the most important fields in today's world, as cyberattacks continue to increase in impact and frequency. Organizations across different sectors face the risks that come with cyber threats such as data breaches, ransomware, phishing attacks and many more, provoking an increasing demand for cybersecurity professionals. According to CyberSeek (CyberSeek), there are more than hundreds of thousands of cybersecurity jobs not filled in the United States, and the need for skilled professionals increases rapidly every year.

To meet this demand, students and early-career cybersecurity professionals pursue cybersecurity certifications and/or competitions as a way to accelerate their chances of securing entry-level cybersecurity jobs. Certifications, such as CompTIA Security+ and Certified Ethical Hacker (CEH), provide detailed knowledge and are most likely to be valued by employers. On the other hand, cybersecurity competitions such as Capture the Flag (CTF), red team/blue team simulations, and other hands-on activities offer practical, real-world experience in threat detection, defense, and exploitation (NCL). Although both cybersecurity certifications and competitions offer valuable learning experiences, they require time, money, and effort. This raises an important question for aspiring cybersecurity professionals: Does having a cybersecurity certification or participating in cybersecurity competitions help individuals secure entry-level positions?

This paper explores that question by analyzing results from a survey of 26 respondents from two key groups: 1) individuals who are currently in entry-level cybersecurity roles, and 2) HR/hiring managers involved in recruitment of cybersecurity staff. The survey contains questions on background, experience, hiring preferences, and recognized value of certifications and competitions. Our findings show that certifications, particularly those from CompTIA and GIAC, are seen as essential, but not sufficient on their own. Participation in competitions is generally viewed as helpful but not required. Our summary result suggests that the most effective path to a cybersecurity role requires a combination of both certifications and real-world experience, and competitions offering helpful value. Future work will include asking for the respondents' business sector since the sector could vary and look at recent trends due to AI usage and adoption.

Keywords: Cybersecurity survey, competitions, certifications, entry-level.

References:

CyberSeek. (n.d.). *CyberSeek*. https://www.cyberseek.org/

National Cyber League (NCL). (n.d.). *National Cyber League*, https://nationalcyberleague.org/

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Cybersecurity capstone project design: A simulation approach

[Poster]

Ping Wang, Robert Morris University, PA, wangp@rmu.edu

Extended Abstract

This poster presents an on-going research and implementation of a simulation approach to designing a comprehensive project for a graduate cybersecurity capstone course. There is a strong industry and workforce demand for cybersecurity professionals to have comprehensive knowledge, skills, and abilities (KSAs) and competencies to address rising and complex cyber threats and risks. Quality graduate education in cybersecurity should develop comprehensive competencies and professional qualifications for problem solving, critical thinking, and leadership for the future of cybersecurity workforce roles and tasks. A capstone project has the potential of providing an integrated learning experience to achieve technical and non-technical learning outcomes and professional competencies for educational assessment and workforce development. The goal of this research is to contribute a proposal with initial empirical implementation data of an enterprise case-based simulation approach to designing a portfolio-type project for the capstone course of a master's degree program in Cybersecurity in the United States.

This research will review the cybersecurity industry and workforce expectations for KSA and competency development, the potential benefits and limitations of the project-based learning (PBL) model, as well as other relevant research on cybersecurity simulation and evaluation. Adopting the case study method with lab simulation, this research uses a selected and sanitized enterprise case organization for identifying and analyzing cybersecurity vulnerabilities, threats, and risks and for designing a comprehensive simulation project for the graduate cybersecurity capstone course. The project simulation design includes both technical and non-technical components and learning activities to develop and evaluate student competencies and KSAs. The OCTAVE (Operationally Critical Threat, Asset, and Vulnerability Evaluation) framework is used for identifying and evaluating the cybersecurity assets, vulnerabilities, and risks for the case organization. The technical simulation features the use of selected SEED Labs for pentesting and hardening for the case study. The proposed project design also includes the key components, objectives, deliverables, and a sample rubric for project assessment and contributes preliminary implementation data on the capstone project design.

Keywords: Cybersecurity, capstone project, project-based learning (PBL), graduate education, simulation, OCTAVE.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

HARM66+ A Taxonomy for the harm-entity-aware post-AI security

[Poster]

Javed Khan, Kent State University, OH, javed@cs.kent.edu

Sharmila Rahman, Kent State University, OH, sprithul@kent.edu

Extended Abstract

The purpose of cybersecurity is harm reduction. Classical cybersecurity has traditionally operated over a narrow and largely fixed set of harms—such as unauthorized access, data loss, service disruption, and financial damage—often framed through extensions of the CIA triad developed for closed, well-bounded systems. While effective in many traditional contexts, these proxy harm models are increasingly strained by emerging system classes. Generative AI systems are open-ended socio-technical actors that generate novel content at scale, influence human cognition and decision-making, and propagate effects across human, institutional, epistemic, and environmental domains. Consequently, they give rise to broader and more complex harms, including loss of human agency, epistemic degradation, erosion of trust, and institutional influence. Many such harms are weakly represented—or absent—in existing cybersecurity taxonomies, allowing systems to satisfy formal security criteria while still producing significant real-world harm. As a result, cybersecurity analysis is constrained when harms cannot be explicitly identified, compared, or reasoned about, and despite extensive prior work, no unified, operationally usable harm taxonomy consistently supports harm analysis across cyber, AI, and socio-technical systems.

In this research, we introduce **HARM66+**, a generated and extensible multi-level harm taxonomy comprising 66+ distinct harm classes organized into two overarching domains and eleven categories. The taxonomy is empirically evaluated for mutual exclusivity, normative orthogonality, completeness, hierarchical traceability, stability, and parsimony. HARM66+ maps harms reported in over 1,500 AI incidents documented by the AI-AAIC Monitor [1] and aligns with thirteen social risk categories derived from an expert review of 1,781 policy, regulatory, and scholarly documents [24]. It was further stress-tested against speculative future harm scenarios drawn from sixteen major science-fiction works. While extensible at lower levels, the upper-level domains and categories are designed for long-term stability and broad applicability.

We now seek community input on HARM66+. A well-designed fine-grained harm taxonomy such as HARM66+ provides foundational benefits to the cybersecurity community. It surfaces i) latent harms and reduces assessment blind spots, ii) enables more harm-specific tools, technology, and process for proportionate, and prioritized monitoring and mitigation strategies, and iii) establishes a shared cross-domain vocabulary for consistent reporting and analysis across institutions, sectors, and jurisdictions. Together, these properties enhance the analytical rigor and practical relevance of cybersecurity research and practice in a post-AI world.

Keywords: Harm taxonomy, cybersecurity foundations, risk scoring, AI and cyber risk governance, resilience analytics, socio-technical security systems.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Applying NIST-based network security controls in a small healthcare clinic

[Poster]

Jose Padilla Carrasquillo, Nova Southeastern University, FL, jp3810@mynsu.nova.edu

Yair Levy, Nova Southeastern University, FL, levyy@nova.edu

Extended Abstract

Healthcare clinics increasingly rely on electronic Protected Health Information (ePHI) to support patient care, billing, and clinical operations. While this reliance improves efficiency, it also expands the cyber-attack surface, particularly for small healthcare organizations with limited security resources. This project examined the cybersecurity posture of a small dental clinic in Florida and proposes targeted network-based security enhancements to mitigate risks to ePHI while maintaining operational feasibility. The clinic's existing environment exhibits critical vulnerabilities, including lack of network segmentation, insufficient access controls, unencrypted internal communications, limited logging, and inadequate device deprovisioning. These weaknesses expose the organization to prevalent healthcare cyber threats such as ransomware, insider misuse, malware propagation, lateral movement, packet sniffing, and denial-of-service attacks. Using a structured risk management approach, this project identified and ranked key cyber risks based on likelihood and organizational impact. The proposed solution aligns with the National Institute of Standards and Technology (NIST) Cybersecurity Framework (CSF) 2.0 and the Cybersecurity McCumber Cube, with adherence to recognized industry best practices.

Recommended controls include the deployment of a centralized intrusion prevention system (IPS), implementation of role-based access control (RBAC), virtual local area network (VLAN) segmentation, encrypted data transmission, centralized logging and monitoring, cloud-based backups, formal device deprovisioning procedures, and user cybersecurity awareness training. These controls collectively address technical, administrative, and operational dimensions of information security. The anticipated outcome of implementing the proposed network architecture is an improved cybersecurity maturity level, elevating the clinic from NIST CSF Tier 1 (Partial) to Tier 2 (Risk-Informed). This improvement reflects enhanced awareness of cyber risks, better integration of cybersecurity practices into daily operations, and stronger resilience against cyber incidents. Additionally, the cost analysis demonstrates that meaningful cybersecurity improvements can be achieved within a modest budget, making the proposal practical for small healthcare providers. Overall, this project illustrates how a risk-informed, standards-based approach can significantly strengthen the protection of ePHI in small clinical environments without imposing excessive financial or operational burdens.

Keywords: Healthcare cybersecurity, electronic protected health information (ePHI), network security, NIST Cybersecurity Framework, intrusion prevention systems, risk management.

References:

National Institute of Standards and Technology. (2024). *The NIST Cybersecurity Framework (CSF) 2.0*. https://doi.org/10.6028/NIST.CSWP.29

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Designing a HIPAA-aligned information security policy program for a small dental clinic

[Poster]

Mauricio Bisogno Loutphi, Nova Southeastern University, FL, mb4383@mynsu.nova.edu

Tomas Francisco Heredia Zuluaga, Nova Southeastern University, FL, th1505@mynsu.nova.edu

Yair Levy, Nova Southeastern University, FL, levyy@nova.edu

Extended Abstract

Small healthcare providers increasingly depend on interconnected digital systems and cloud-based applications to manage electronic protected health information (ePHI). While these technologies improve efficiency and patient care, they also introduce significant cybersecurity and regulatory compliance risks—particularly for small dental clinics that often lack formal governance structures and dedicated security resources. This poster presents a policy-driven program to managing ePHI risk through the development of an information security policy framework and compliance plan aligned with the Health Insurance Portability and Accountability Act (HIPAA) Security Rule and the National Institute of Standards and Technology (NIST) Cybersecurity Framework (CSF) 2.0.

A case-based assessment of a small, single-site dental clinic revealed multiple high-risk gaps, including shared user credentials, absence of multifactor authentication, inconsistent encryption of ePHI, limited cybersecurity awareness training, weak audit logging and monitoring, insufficient oversight of third-party vendors, and inadequate physical and media controls. These weaknesses expose the clinic to common healthcare threats such as phishing, ransomware, insider misuse, and supply-chain compromise, while also increasing the likelihood of noncompliance with HIPAA administrative, technical, and physical safeguards.

To address these risks, the proposed solution focused on governance through seven interrelated information security policies: (1) Access Control and Authentication, (2) Cybersecurity Awareness and Training, (3) Data Protection and Encryption, (4) Audit Logging and Monitoring, (5) Third-Party and Remote Access Management, (6) Device and Media Controls, and (7) Physical Security and Facility Access Control. Each policy is explicitly mapped to relevant HIPAA Security Rule provisions and NIST CSF 2.0 outcomes within the Identify, Protect, and Detect functions, enabling regulatory traceability and audit readiness. The anticipated outcome of implementing this program is a transition from an ad hoc, person-dependent cybersecurity posture to a documented, repeatable, and risk-informed program. By emphasizing policy development, workforce accountability, and scalable controls, this program demonstrates how small healthcare practices can meaningfully reduce cybersecurity risk, strengthen protection of ePHI, and support operational resilience without excessive cost or complexity.

Keywords: Healthcare cybersecurity, HIPAA Security Rule, NIST Cybersecurity Framework, information security policy, ePHI protection, risk management.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Securing the legacy: Aspect-oriented forward engineering of OCL constraints in brownfield development

[Poster]

Ishi Golub, Wentworth Institute of Technology, MA, golubi@wit.edu

Aspen Olmsted, Wentworth Institute of Technology, MA, olmsteda@wit.edu

Extended Abstract

Modern software development frequently involves "brownfield" projects where new features must be integrated into large, complex legacy codebases. As development cycles accelerate with AI-assisted programming, the risk of violating critical business rules and security invariants in these legacy systems increases. This poster presents a methodology for securing brownfield development by forward engineering Object Constraint Language (OCL) constraints from formal models into executable enforcement layers using Aspect-Oriented Programming (AOP).

Methodology - Our approach bridges the gap between high-level architectural design and legacy implementation without requiring a complete rewrite of the underlying codebase. The workflow follows three primary phases: We identify critical security and domain invariants and formalize them using OCL within a UML model. Rather than manually embedding check logic into the legacy source code—which is error-prone and invasive—we forward-engineer these OCL constraints into Aspects. Using AOP, these constraints are "woven" into the system at specific join points. This approach allows the security logic to remain modular and external to the core legacy logic, ensuring that any code generated by AI or written by developers is automatically validated against formal constraints at runtime.

Experimental Results - The methodology was validated using SuiteCRM, a popular open-source PHP-based CRM that represents a typical complex brownfield environment, and we found that applying AOP-driven OCL enforcement successfully intercepted unauthorized data operations that bypassed standard UI-level checks. The overhead of the woven security aspects remained negligible, while the system's security posture was significantly hardened. The approach can be applied in any software project where the appropriate modelling assets exists and AOP is supported by the programming language.

Conclusion - This research demonstrates that formal modeling and AOP provide a robust framework for retrofitting security into legacy systems. By treating OCL constraints as first-class, injectable citizens, developers can safely evolve brownfield applications such as SuiteCRM, ensuring that modern development speed does not compromise historical system integrity.

Keywords: Brownfield Development, Object Constraint Language (OCL), Aspect-Oriented Programming (AOP), formal modeling, software security invariants.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Federated learning-based anomaly detection for high-performance research networks

[Poster]

Ehsan Saeedizade, University of Nevada, Reno, NV, esaeedizade@unr.edu

Shamik Sengupta, University of Nevada, Reno, NV, ssengupta@unr.edu

Jay Thom, University of Nevada, Reno, NV, jthom@unr.edu

Extended Abstract

High-performance research networks (HPRNs) underpin data-intensive scientific workflows, yet they are increasingly exposed to network anomalies and cyber threats such as operational faults, misconfigurations, and adversarial activities that can degrade performance, compromise security, reliability, and availability, and undermine trust in distributed cyber-infrastructure. Detecting such anomalies in a timely and automated manner is an important task yet challenging due to the volume of telemetry and privacy constraints across administrative domains that prevent sharing the telemetry data. Traditional rule-based and centralized machine-learning (ML) anomaly detection methods depend on centralized training, which not only poses privacy risks but also encounters challenges related to scalability and limited generalization across heterogeneous networks with non-IID behaviors.

To address these issues, this poster presents a privacy-preserving, federated learning-based anomaly detection approach designed for cross-site monitoring of HPRNs. In this approach, each participating site trains a local model using fine-grained host, network, and storage metrics, and only model updates are shared with a coordinating server, avoiding the centralization of sensitive operational data. To address domain shift and non-IID data across sites, it integrates a feature-based transfer learning mechanism that improves model generalization across non-IID datasets and enhances robustness in environments with sparse data, thus mitigating the adverse effects of heterogeneous infrastructure conditions and limited training data availability. For example, when research facilities or universities perform large-scale scientific data transfers between their sites, each site monitors local infrastructure and collects telemetry during transfers.

The participating sites collaboratively train an ML model using monitored telemetry and following the proposed approach. Once the model is trained, each site deploys the learned model locally, continuously monitors transfers, and automatically detects anomalous behavior in real time. The proposed approach is extensively evaluated on datasets generated from various testbeds and achieves a high detection accuracy of 96.56% and an F1-score of 0.9658, which significantly outperforms centralized training paradigms. Moreover, experimental results show that feature normalization increases the convergence of federated learning model training, reducing the number of rounds required for stabilization from 100 to 50 rounds. These findings highlight the efficacy of our proposed federated learning-based approach in securing HPRNs and offering a scalable, privacy-preserving solution for anomaly detection in distributed research infrastructures.

Keywords: Federated learning, anomaly detection, data privacy, file transfer, transfer learning.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Toward AI-driven monitoring and malware detection framework for IoT and edge systems

[Poster]

Ehsan Saeedizade, University of Nevada, Reno, NV, esaeedizade@unr.edu

Jake Lyon, The College of Wooster, OH, jlyon26@wooster.edu

Shamik Sengupta, University of Nevada, Reno, NV, ssengupta@unr.edu

Extended Abstract

The rapid expansion of Internet of Things (IoT) and edge devices across smart cities, transportation, and industrial systems has significantly increased the attack surface of modern cyber–physical infrastructures. These devices are often resource-constrained, physically exposed, and deployed in highly heterogeneous network environments, making them attractive targets for malware, botnets, and coordinated cyberattacks. Traditional SOC and monitoring frameworks were primarily designed for centralized enterprise networks and struggle to operate effectively in large-scale IoT and edge environments due to limited device visibility, high telemetry volumes with heterogeneous data formats, fragmented data sources, and manual analysis of data by security analysts, thereby not suitable for real-time analysis at scale.

As a result, next-generation SOC platforms are increasingly adopting AI-enabled monitoring and detection pipelines to automate threat detection and reduce operational burden. Machine learning (ML) offers a promising approach for automated malware detection and classification; however, practical deployment in IoT environments requires models that are not only accurate but also lightweight, data-efficient, and robust to evolving threats. In this work, we present an AI-driven monitoring and malware detection framework designed for IoT and edge environments and investigate the effectiveness of supervised ML models for malware detection tasks. The framework integrates continuous edge telemetry collection with centralized analytics to support SOC workflows. We then investigate the effectiveness of four supervised learning models (Random Forest, LightGBM, Logistic Regression, and a Multi-Layer Perceptron) across binary malware detection and multiclass malware family classification. We further assess sensitivity to limited labeled data and temporal robustness to simulate deployment under an evolving cyber threat landscape and concept drift using the IoT-23 dataset.

Our results show that tree-based models achieve strong accuracy and generalization even with small training sets, while all models exhibit performance degradation over time as malware diversity increases. These findings emphasize the importance of adaptive, resource-efficient, and continuous model training approaches for improving security in real-world IoT and edge deployments.

Keywords: Malware detection, Internet of Things, machine learning, cybersecurity, AI-driven SOC, monitoring.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Modeling cloud-controlled cyber-physical system resilience using a reinforcement-learning cart-pole testbed

[Poster]

Imran Jahid Khan, Tuskegee University, AL, mkhan7037@tuskegee.edu

Mohammad Rahman, Tuskegee University, AL, mrahman@tuskegee.edu

Fan Wu, Tuskegee University, AL, fwu@tuskegee.edu

Extended Abstract

Modern infrastructures increasingly depend on smart software and learning-enabled controllers to manage physical machines. Although cloud-based control increases operational effectiveness and scalability, it also makes these systems vulnerable to cyberattacks. Thus, system resilience under adversarial conditions is a key concern in cyber defense contexts.

This poster presents a simulated cart-pole system modeled as a cloud-controlled cyber-physical system. The cloud-based control component is represented by a reinforcement-learning controller, which communicates with the physical system through sensor and actuator feedback loops. The testbed is designed to investigate how the system behaves under both horizontal and inclined surface conditions, representing normal and stressed operating environments. Observation noise and other cyber-relevant threats such as sensor spoofing, signal interference, and data integrity degradation are incorporated into the model. These setups enable a systematic investigation of how cyber interference could spread throughout the control loop and affect system stability, recovery behavior, and overall assurance.

Intuitive performance indicators, including survival duration, failure frequency, and recovery behavior, are used to assess resilience in cloud-controlled cyber-physical systems. These metrics are intended to support analysis of how inclined operating conditions and adversarial inputs can affect system stability under environmental stress. The use of horizontal and inclined surface conditions enables the testbed to represent a range of operating scenarios commonly encountered in real-world cyber-physical systems, including robotic platforms and space-related applications. As a result, the suggested framework offers a versatile and understandable foundation for studying how environmental stress and cyber interference interact in learning-enabled, cloud-based cyber-physical systems.

Keywords: Cyber-physical systems, cloud-controlled systems, cyber resilience, reinforcement learning, adversarial interference, cart-pole testbed.

References:

Yu, Z., Gao, H., Cong, X., Wu, N., & Song, H. H. (2023). *A survey on cyber–physical systems security*. *IEEE Internet of Things Journal, 10*(24), 21670–21686. https://doi.org/10.1109/JIOT.2023.3289625

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Building a privacy and security layer around LLM models to protect against common AI attacks

[Poster]

Samuel Kadima, Xavier University, OH, kadimas@xavier.edu

Kyle Totorica, Xavier University, OH, totoricak@xavier.edu

Extended Abstract

Large Language Models (LLMs), such ChatGPT, Copilot, etc., have become integral parts of our daily lives, assisting us with learning, productivity, and various tasks. These AI systems are valuable because they help with tedious tasks. However, LLMs are stateless by default and lack built-in mechanisms for memory and security; as a result, threat actors can easily manipulate them to reveal private or sensitive information, which may result in unsafe, harmful, or misleading outputs because LLMs respond to user queries without understanding the intent.

There are several common AI attacks against LLMs that raise serious concerns. For example, prompt injection occurs when a threat actor manipulates the LLM to generate malicious output by following their instructions or overriding the system rules, leading to unintended behavior. Additionally, attackers could use obfuscation techniques like Base64 encoding or other forms to hide malicious inputs.

LLMs also struggle to protect sensitive information, including personally identifiable information (PII), such as financial data and health records, which can lead to unintended exposure of confidential information without proper guardrails. In addition, a lack of input and output handling could allow LLMs to produce unsafe outputs to user prompts without proper validation or sanitization. These challenges demonstrate the importance of having a security layer around LLM models.

This project highlights the importance of AI security and proposes a privacy and security layer based on zero-trust principles (a safety shield for LLMs), where neither the user input nor the model outputs are trusted by default. An orchestrator acts as the central component and policy enforcement, controlling the interaction between the user and the LLM. This works by implementing security checks on user inputs before they reach the model and reviewing the model outputs before sending them back to the user.

Keywords: Artificial intelligence (AI), AI security, large language models, safe AI use, prompt injection, insecure output handling, sensitive information disclosure.

Jericho: An accessible cyber city model for teaching cyber operations

[Poster]

Alex Taylor, Cedarville University, OH, alexandertaylor@cedarville.edu

Andrew Quick, Cedarville University, OH, aquick@cedarville.edu

David Moore, Cedarville University, OH, davidmoore155@cedarville.edu

Nathaniel Gavilan, Cedarville University, OH, ngavilan@cedarville.edu

Seth Hamman, Cedarville University, OH, shamman@cedarville.edu

Extended Abstract

Traditionally, educational institutions have conducted cybersecurity exercises exclusively in virtual environments, reducing student awareness of the profound impacts cyberattacks can have in physical space and on critical infrastructure. The Jericho project targets students from secondary education to advanced undergraduate level. Jericho drives home the potential physical space effects of cyberattacks by blending virtual computer networks with a physical city model. Other cyber-physical sandbox environments exist, but they are prohibitively expensive and found only at a few select institutions. Jericho is an affordable and scalable cyber city built using inexpensive wooden crates, Raspberry Pis, 3D printed structures, and low-cost model railroad artifacts. Additionally, Jericho's hacking scenarios are accessed through an intuitive web application ensuring that students of all ages and technical experience can master core cybersecurity competencies. Jericho's web application includes a Kali Linux virtual machine web console, a question-and-answer scoring system, and a livestream of the physical city fed by strategically placed cameras. With these features, students can complete custom-designed scenarios while simultaneously observing the effects their actions have on the physical infrastructure. Scenarios can be spun up and torn down on demand in a traditional cyber range environment that is ultimately networked with Raspberry Pis in the city capable of creating physical effects.

Keywords: Accessible, affordable, critical infrastructure, cyber city, cyberattacks, web application.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Securing AI systems through LLM red teaming: A practical pillar of modern AI governance

[Poster]

Kellep A. Charles, Capitol Technology University, MD, kacharles@captechu.edu

Extended Abstract

As large language models and other advanced AI systems move rapidly from experimentation into mission-critical business functions, traditional security and compliance controls are proving insufficient to manage their unique risk profile. Unlike conventional software, AI systems can generate unpredictable outputs, amplify latent bias, leak sensitive information, and be manipulated through adversarial prompts or data poisoning. These risks introduce new attack surfaces that sit at the intersection of cybersecurity, ethics, and organizational governance.

This session examines large language model red teaming as a core security practice within a mature AI governance program. Red teaming goes beyond model accuracy testing to systematically probe AI systems for vulnerabilities related to confidentiality, integrity, availability, safety, and trust. Participants will explore how structured adversarial testing can uncover issues such as prompt injection, jailbreaks, training data leakage, hallucinations, harmful content generation, and failure modes that emerge under real-world conditions.

To illustrate its practical application, consider a healthcare organization deploying an AI-powered clinical support assistant. Through red teaming, testers simulate adversarial prompts designed to extract sensitive patient information, induce unsafe medical recommendations, or bypass safety guardrails. The results reveal gaps in data handling, model alignment, and output validation, which are then fed back into governance processes to strengthen controls, refine policies, and reduce risk prior to full deployment.

The discussion positions LLM red teaming not as a one-time technical exercise, but as a continuous governance control aligned with enterprise risk management. It outlines how red teaming activities can be operationalized across the AI lifecycle, from pre-deployment assessments and model selection to post-deployment monitoring and incident response. The session also highlights how red teaming outputs feed into AI risk registers, impact assessments, model documentation, and executive reporting, enabling defensible decision-making for leadership and regulators.

By framing LLM red teaming as both a security discipline and a governance mechanism, this session provides practitioners, risk leaders, and executives with a practical approach to strengthening AI resilience, improving transparency, and demonstrating responsible stewardship of AI systems in an increasingly regulated and adversarial landscape.

Keywords: AI red teaming, AI governance, Large Language Models (LLMs), adversarial machine learning, AI risk management.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Free and open-source Security Orchestration, Automation, and Response Platform (FOSS-SOAR)

[Poster]

Aedan Podest, Xavier University, OH, podesta@xavier.edu

Cooper Landen, Xavier University, OH, landenc1@xavier.edu

Extended Abstract

Security Operations Centers (SOCs) increasingly rely on Security Orchestration, Automation, and Response (SOAR) platforms to streamline alert triage, enrichment, and response. While commercial SOAR solutions are common in industry, their proprietary nature and licensing requirements limit accessibility for individual learners seeking academic exploration within incident response. As a result, students and early-career practitioners often have limited opportunities to examine how automated security workflows are designed and executed from end to end.

This project presents FOSS-SOAR, a free and open-source security orchestration, automation, and response platform developed as a senior cybersecurity capstone project. The goal was to design and implement open-source tools to explore the structure, functionality, and tradeoffs of security automation within incident response. The project focuses on building a functional SOAR platform within a virtual environment using freely available tools, with an emphasis on transparency and hands-on experimentation relevant to cyber defense education. Rather than serving as a production-ready system, FOSS-SOAR is intended as proof-of-concept that demonstrates how common SOC tasks can be automated and orchestrated within a controlled environment.

FOSS-SOAR ingests events generated from simulated activity and open-source security tools including Wazuh for centralized logging, Suricata and Sysmon for network and host telemetry, and integrates alerts with TheHive and Cortex for case management and automated analysis. These events are processed through configurable playbooks that perform enrichment, apply conditional logic, and execute response actions such as logging, notifications, or containment. All automation logic is implemented using open-source components, allowing the internal decision-making process of each playbook to remain visible and inspectable. A key objective of the project is to examine the balance between automation and human judgment from an incident response perspective. While FOSS-SOAR automates repetitive steps such as data enrichment and alert triage, it is intentionally designed to avoid fully autonomous decision-making, requiring analyst interaction at key decision points in the workflow.

Keywords: Incident response, open-source, automation, SOC workflow, SOAR.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Use of Cyber-Informed Engineering for digital risk mitigation

[Poster]

Wm. Arthur Conklin, University of Houston, TX, waconklin@uh.edu

Extended Abstract

This poster presents a series of graphical elements to assist in the communication of the role Cyber-Informed Engineering (CIE) process plays in assisting engineers in the addressing of digital risks that are becoming commonplace in engineered systems as computers add intelligence to modern systems. The poster will emphasize how engineered systems have different security objectives than IT systems and how this results in a different approach to security. Safety, resilience and operational efficiency replace the CIA triad, and are specifically designed into a system. The CIE program has 12 principles which are listed and a new construct called an engineered control to achieve desired level of protection. An examination of the types of engineered controls and how they enable the digital risk mitigations necessary in a system is presented. The poster will also have references in the form of QR codes to documents from the CIE program to assist in the navigation of the details needed to include elements in a CAE program.

The primary purpose of this unattended poster is to act as a communication mechanism to introduce the topic in a series of short elevator pitches and links to relevant documentation developed as part of this ongoing effort. Key documents include CIE Implementation Guide, A curriculum guide and curriculum resources, CIE CMM materials, and Engineered Controls documentation.

Posters provide an efficient means of communicating to members through a series of graphical elements backed by URL links for deeper dives.

Keywords: Digital risk mitigation, Engineered controls, Cyber-Informed Engineering (CIE).

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

US Coast Guard Academy (USCGA) eCTF 2026

[Poster]

Tyler Bissett, US Coast Guard Academy, USA, tyler.m.bissett@uscga.edu

Phillip Kuznetsov, US Coast Guard Academy, USA, phillip.v.kuznetsov@uscga.edu

Caden Waters, US Coast Guard Academy, USA, caden.l.waters@uscga.edu

Daniel Yi, US Coast Guard Academy, USA, daniel.y.yi@uscga.edu

Mohamed Elwakil, US Coast Guard Academy, USA, mohamed.m.elwakil@uscga.edu

Extended Abstract

This poster introduces a hardened, permission-based file storage and transfer system designed for the 2026 MITRE Embedded Capture the Flag (eCTF) competition. The system is engineered to protect sensitive intellectual property—specifically semiconductor design files, firmware updates, and calibration data—within a simulated manufacturing environment. The architecture is centered around three distinct Hardware Security Modules (HSMs): the Engineer HSM for creating and managing proprietary designs, the Technician HSM for system maintenance and updates, and the Photolithography Machine HSM, which consumes secure files. These HSMs, built on the MSPM0L2228 microcontroller, communicate via a secure UART channel, all managed by a host computer, to prevent unauthorized access and data exfiltration by a malicious insider.

To meet the competition's rigorous security challenges, the system implements a multi-layered defense strategy. Functionally, it relies on a secure build process that generates deployment-specific secrets and embeds a strict, group-based permission model (Read, Write, Receive) into each HSM's firmware. Security is enforced through several key mechanisms: all sensitive file operations are protected by a secret PIN to thwart attacks even with physical device access; a custom authenticated and encrypted protocol secures all inter-HSM file transfers to guarantee data integrity and confidentiality against man-in-the-middle attacks; and cryptographic hashes are employed to verify the authenticity of all firmware and design files, preventing malicious modification, corruption, or the introduction of backdoors.

Keywords: Hardware security module, MITRE eCTF 2026, secure flash storage, Authenticated Encryption (AEAD), permission-based access control.

References:

MITRE. (n.d.). *Embedded capture the flag (eCTF) competition.* https://ectfmitre.gitlab.io/ectf-website/index.html

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Building a secure, scalable capture the flag platform for cybersecurity education

[Poster]

Giovanni Braun, University of Southern Maine, ME, giovanni.braun@maine.edu

Extended Abstract

The purpose of this project was to design and deploy a custom Capture The Flag environment that could serve as both an instructional tool and a reusable institutional resource. The system was intended to expose participants to multiple domains of cybersecurity, including cryptography, networking, open-source intelligence, password cracking, and reverse engineering. A core requirement of the project was reproducibility, ensuring that future instructors or administrators could deploy, modify, and maintain the environment with minimal effort. The environment was built using modern infrastructure practices such as containerization, orchestration, and secure network access. These technologies enabled the creation of a CTF platform capable of supporting diverse challenges, multiple users, and realistic cybersecurity scenarios in a controlled and secure manner.

The Capture The Flag competition was structured around five primary cybersecurity domains: cryptography, networking, open-source intelligence (OSINT), password cracking, and reverse engineering. These categories provided students with broad exposure to key areas of cybersecurity while allowing them to engage with both analytical and technical problem-solving tasks.

The environment was hosted on an Ubuntu Server system, which served as the primary infrastructure host. The CTFd platform was deployed within a Docker container, providing isolation from the host operating system and ensuring consistent behavior across deployments.

The challenges were developed to align with educational objectives while providing realistic cybersecurity scenarios. Participants were required to analyze and break common encryption techniques, analyze packet captures along with live traffic interception with tools such as Wireshark and tcpdump. Reverse engineering challenges required students to inspect binaries, analyze program logic, and extract embedded secrets. Reverse engineering consistently proved to be the most difficult category, highlighting its complexity and the need for strong foundational knowledge. Open-source intelligence challenges were based on real-world data. Students applied OSINT techniques such as geolocation, metadata analysis, and open-source research.

Secure, scalable, and reproducible Capture The Flag environments can be successfully deployed using open-source tools and modern infrastructure practices. By combining containerization, orchestration, and secure networking, the system provides a robust platform for cybersecurity education. The design enables future instructors to recreate and extend the environment to support use in academic and training settings.

Keywords: CTF, hands-on learning, CTF design methodology, scalable training environment, cybersecurity education.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

From mass extraction to ethical triage: Multi-agent AI for privacy-preserving computer forensics

[Poster]

Md Tamjid Hossain, Texas A&M University-San Antonio, TX, mhossain@tamusa.edu

Shafkat Islam, Purdue University Northwest, IN, islam59@pnw.edu

Extended Abstract

A central ethical dilemma in computer forensic investigations stems from a very practical tension: investigators want to collect *everything* to avoid missing critical evidence, yet doing so inevitably exposes vast amounts of personal and sensitive data that have nothing to do with the case at hand. Contemporary investigative practices increasingly rely on large-scale data collection from heterogeneous sources, including disk images, mobile devices, cloud logs, and network traces, where relevant artifacts are deeply intermingled with private communications, personal photos, medical records, and other nonresponsive information. Traditional "whole-disk acquisition" or "mass extraction" approaches may be operationally efficient, but they force forensic investigators to sift through far more data than the investigative scope justifies.

This reality raises serious concerns around privacy overreach, legal compliance, and ethical responsibility, particularly under frameworks such as GDPR, CCPA, HIPAA, and emerging privacy engineering guidelines. Our judicial systems have also recognized that doctrines developed for physical searches do not translate cleanly to digital environments, where irrelevant and relevant data are unavoidably co-located. In response, courts have begun imposing stricter limits during the search authorization phase. However, overly restrictive collection policies introduce a different risk: excluding exculpatory or contextual artifacts that may be essential for reconstructing events accurately. As a result, investigators are often left navigating a difficult trade-off between protecting privacy and preserving investigative completeness- without adequate technical support to do both well.

This work argues that ethical cybercrime investigation practice must move beyond reactive, tool-centric workflows and toward *ethics-by-design* investigative systems. Building on our ongoing research, we propose a privacy-preserving approach that embeds *multi-agent AI reasoning* directly into the forensic pipeline. Rather than treating all acquired data as equally relevant, a coordinated set of AI agents analyzes forensic artifacts in context, using case narratives, timelines, and behavioral indicators to assess whether each artifact is likely to be case-related. Investigators see only what is likely relevant by default, while sealed data remains accessible under explicit authorization and policy oversight. Thus, from an ethical standpoint, this approach directly addresses two core privacy concerns: (a) data minimization, and (b) accountability. Within the broader CAE-CD mission, these concerns are not solely technical but educational and workforce-oriented. Future cybersecurity professionals must be trained to reason about privacy, ethics, and legal boundaries as first-class design constraints, not afterthoughts.

Keywords: Ethical computer forensics investigations, multi-agent AI, data minimization, policy-regulated access control.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Evansdale 2050: A cybersecurity-centered cyber-physical model of a Personal Rapid Transit system

[Poster]

Tucker Amon, West Virginia University, Morgantown, WV

Salem Hefeida, West Virginia University, Morgantown, WV

Jordan Luzier, West Virginia University, Morgantown, WV

James Sleptzoff, West Virginia University, Morgantown, WV

Cosmas Nwakanma, West Virginia University, Morgantown, WV

Mohamed Hefeida, West Virginia University, Morgantown, WV, mohamed.hefeida@mail.wvu.edu

Extended Abstract

Personal Rapid Transit (PRT) systems constitute safety-critical cyber-physical systems that rely on tightly coupled industrial control, real-time networking, and supervisory software. As such systems evolve toward greater autonomy and connectivity, they inherit many of the vulnerabilities observed in industrial control systems (ICS) and supervisory control and data acquisition (SCADA) environments. This work presents a cybersecurity-centered modeling and experimental evaluation of the West Virginia University (WVU) PRT system, with emphasis on identifying, exploiting, and mitigating vulnerabilities within an operationally realistic cyber-physical environment (Evansdale 2050 Testbed). The proposed framework integrates a physical scale model of the Evansdale Campus with programmable logic controllers (PLCs), sensors, and a desktop-based supervisory control environment, complemented by simulation tools to emulate system-wide behavior. Real-world operational factors—including class schedules, peak travel periods, and event-driven congestion—are incorporated to model daily demand patterns and stress-test system performance. Particular attention is given to translating simplified physical mechanisms, such as conveyor-based movement, into accurate representations of independently operating PRT vehicles. A core contribution of this work is a layered cybersecurity assessment of the PRT architecture. The study systematically evaluates vulnerabilities across front-end interfaces, back-end control systems, and PLC components, examining threats such as injection attacks, authentication weaknesses, denial-of-service and replay attacks, and unauthorized data access. Network segmentation and firewall-based access control are implemented to compare "secure" and "vulnerable" system states, enabling controlled experimentation and attack–defense analysis. Expected outcomes include a hardened and extensible PRT model, improved documentation for reproducibility, and alternative control and congestion-management algorithms beyond traditional request-and-report mechanisms. This research advances the understanding of secure cyber-physical transportation systems and provides design insights applicable to future smart-city-scale PRT deployments.

Keywords: Cyber-Physical Systems, Transportation Cybersecurity, Industrial Control Systems, PLC Security, Smart Cities, Personal Rapid Transit.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Securing real-time biosignal command streaming for assistive robotics using AES-GCM

[Poster]

Joy Williams, Morgan State University, MD, jowil26@morgan.edu

Chukwuebuke Okwese, Morgan State University, MD, chokw9@morgan.edu

Onyema Osuagwu, Morgan State University, MD, onyema.osuagwu@morgan.edu

Extended Abstract

Biotechnology control pipelines are being increasingly used in assistive and rehabilitation therapy, but transmitting physiological data introduces major cybersecurity risks. Electrooculography (EOG) can be mapped to directional commands (e.g., left/right/up/down), creating a direct pathway from human physiology to cyber-physical execution. Without protection, adversaries may attack by injecting, replaying, tampering with, or flooding traffic, potentially producing unsafe or unintended behavior.

This poster presents a secure, real-time EOG-to-robot command pipeline implemented in a simulation. EOG signals are received from an Open BCI Galea, a wearable Brain Computer Interface, and streamed into a VR Unity application. Guided calibration steps establish a baseline for directional references, then horizontal and vertical EOG components are mapped into continuous control values. These control values are transmitted via UDP to a Python-based receiver that interfaces with a MuJoCo simulation of the Unitree G1 robot. Our setup enables quick prototyping of hands-free control while supporting security evaluation in a safe environment, providing a baseline for future work on ocular motor control and neural signal security in assistive and neurodegenerative research.

To address network-level threats, the system applies AES-GCM authenticated encryption of transmitted control packets. Each packet includes a unique identifier and a timestamp. The receiver can reject malicious commands that have been tampered with, replayed, or stale. An experimental evaluation was conducted using packet injection/tampering, replay attacks, and traffic flooding scenarios. Consistently, the protected channel rejected thousands of malicious packets through authentication failures, replay detection, and timestamp validation. Running in parallel, the unsecured channel accepted all traffic without verification. The secured pipeline sustained average command latencies of 52-77ms while parameters were varied for each attack. Comparable to the unsecured channel, AES-GCM and replay checks introduced minimal overhead.

The main contribution of this work to the CAE-CD community is a reproducible case study that demonstrates how standard encryption and authentication can be integrated into real-time physiological data transmission without refactoring the sensing control stack and disrupting responsiveness. These tests measure behavioral response to malicious activity and highlight security-latency tradeoffs. This is relevant to assistive robotics and healthcare-adjacent cyber-physical systems.

Keywords: Electrooculography, assistive robotics, cyber-physical systems, AES-GCM, authenticated encryption, secure neural data.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

SEMANTIC search for healthcare patient data using sentence transformers and ChromaDB

[Poster]

Megan Rabb, South Carolina State University, SC, mrabb1@bulldogs.onmicrosoft.com

Shani Walker, South Carolina State University, SC, swalke18@bulldogs.onmicrosoft.com

Joniqua Bates, South Carolina State University, SC, jbates9@bulldogs.onmicrosoft.com

Xiaomao Liu, South Carolina State University, SC, xliu@scsu.edu

Janmejay Mohanty, South Carolina State University, SC, jmohanty@scsu.edu

Biswajit Biswal, South Carolina State University, SC, bbiswaji@scsu.edu

Nikunja Swain, South Carolina State University, SC, swain@scsu.edu

Extended Abstract

Healthcare systems often store large volumes of patient records in formats that are difficult to search efficiently using traditional keyword-based methods. These limitations can delay care, frustrate providers, and impact outcomes. A solution is needed that understands the context and meaning of clinical documentation, enabling faster, more intelligent access to related cases.

Semantic search changes how we manage healthcare information by helping us understand the meaning behind patient records rather than just looking for exact words. In this project, we used Google Colab to build a simple but powerful semantic search system that combines Sentence Transformers and ChromaDB. The goal was to make it easier to find similar patient cases or notes based on meaning. We used a pre-trained transformer model ("all-MiniLM-L6-v2") to turn sentences about patient data into numerical vectors. These vectors were saved and searched using ChromaDB, a lightweight vector database. All coding and testing were done in Google Colab. For example, when we searched for "high blood pressure treatment," the system returned a sentence about "medication for hypertension" proving that it could understand medical terms even if they were worded differently. This kind of system can make it easier for doctors or medical staff to quickly find relevant records, especially in electronic health record systems. Overall, this project demonstrates how machine learning tools, such as sentence embeddings, can make healthcare data more intelligent and useful.

Keywords: ChromaDB, semantic search engine, Google Colab, patient records, sentence transformers.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Development of unified theoretical framework for zero trust enterprise network cyber security

[Poster]

Olufemi Agunbiade, Morgan State University, MD, olagu4@morgan.edu

Onyema Osuagwu, Morgan State University, MD, onyema.osuagwu@morgan.edu

Extended Abstract

Zero Trust Cybersecurity (ZTCS) represents a fundamental paradigm shift toward continuous verification and eliminating implicit trust. While widely adopted across industry and government, the field lacks a cohesive theoretical framework for implementing, measuring, and comparing Zero Trust Architecture (ZTA) effectiveness. This research addresses this gap by developing a unified mathematical framework for Zero-Trust Enterprise Network Security (ZTENS) using measurement-driven theory and quantitative method to evaluate an enterprise network ZT level.

The framework models access control decisions as repeated statistical evaluations under uncertainty, where each session receives a dynamic trust score derived from multiple evidence signals: identity authentication, device posture, behavioral analytics, network integrity, environmental context, and policy compliance. These signals are probabilistically fused using Bayesian posterior theory, enabling continuous trust estimation that updates as new verification evidence arrives. Privilege elevation occurs only when sufficient information gain justifies additional authentication steps, balanced against operational friction and mission utility.

A critical theoretical insight is that absolute zero-risk Zero Trust is unachievable in practical systems due to fundamental uncertainty and computational limits. Instead, the research introduces the Minimal-Trust System (MTS) concept, where implicit trust is systematically minimized through continuous verification and calibrated risk thresholds. Risk reduction is achieved by increasing the quantity, diversity, and freshness of independent verification signals.

To enable organizational benchmarking, the research proposes the Zero Trust Maturity Index (ZTMI)—a quantitative metric measuring how closely enterprise architectures approach the minimal-trust boundary. Trust scores of 10 enterprise network ranged from 0.35 to 0.88, indicating maturity levels from partial to near-complete Zero Trust adoption. Results show that strong authentication mechanisms, frequent telemetry collection, and automated policy enforcement significantly improve Zero Trust maturity, while weak device posture reduces scores.

Beyond theory, this framework advances cybersecurity education by providing quantitative foundations for teaching Zero Trust, enabling graduate-level laboratory simulations of trust-score computation, verification policies, and maturity assessment using real enterprise telemetry.

Keywords: Zero Trust Architecture, enterprise network security, trust metrics, quantitative evaluation, security maturity models, cybersecurity education.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Development of a digital forensics lab for incident response, criminal investigation, and host analysis training

[Poster]

William Meredith, West Virginia University, WV, wm00026@mix.wvu.edu

Thomas R Devine, West Virginia University, WV, thomas.devine@mail.wvu.edu

Extended Abstract

The growing demand for skilled cybersecurity professionals has elevated digital forensics to a core component of cyber defense education. Digital forensics skills are essential for tasks such as incident response, host-based compromise analysis, and criminal investigation. However, many academic programs struggle to provide students with realistic, hands-on forensic experiences that mirror professional practice. This poster presents the design and classroom use of a modular digital forensics lab framework developed for undergraduate cybersecurity education.

The lab framework is designed around realistic, scenario-based investigations using prepared disk images and open-source forensic tools. The instructional focus emphasizes forensic process and analytical reasoning rather than step-by-step tool usage. The lab aligns with established digital forensics principles, including evidence identification, artifact analysis, timeline reconstruction, and professional reporting, and is intended to be reusable and adaptable across cyber defense curricula.

The lab consists of two independent investigative scenarios. The first scenario uses a removable flash drive image containing deleted documents and images. Students recover deleted files, analyze file system artifacts, and reconstruct user activity to determine what data was present and how it was altered. The second scenario simulates a Linux-based enterprise system compromise in which an attacker escalates privileges and deletes authorization and credential logs to obscure activity. Students examine host-based artifacts to identify indicators of compromise, reconstruct attacker actions, and assess system impact.

Student learning is supported through required documentation of findings and concise forensic reporting that reflects professional practice. To assess the educational utility of the exercise, students complete a post-lab survey using standardized questions to evaluate perceived learning gains and confidence in applying digital forensics techniques.

The lab supports multiple cyber defense workforce roles defined by the Department of Defense Cyber Workforce Framework, including Cyber Defense Forensics Analyst and Cyber Crime Investigator. Initial classroom deployments suggest strong student engagement and perceived value in applying forensic techniques to realistic investigative scenarios.

This work contributes a practical digital forensics lab framework that balances realism, instructional accessibility, and ease of adoption. Future work includes refining scenarios, expanding survey-based assessment, and releasing the lab materials to the CAE community to support shared curriculum development in cyber defense education.

Keywords: Digital Forensics, Cybersecurity Education, Incident Response, Experiential Learning, Cyber Defense Workforce.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

The impact of scholarships and co-curriculum activities on the academic success and career prospects of cybersecurity students

[Poster]

Katerina Goseva-Popstojanova, West Virginia University, WV, katerina.goseva@mail.wvu.edu

Robin Hensel, West Virginia University, WV, robin.hensel@mail.wvu.edu

Extended Abstract

To contribute towards addressing the national shortage of cybersecurity experts, a new B.S. degree and an Area of Emphasis (AoE) in Cybersecurity were developed at West Virginia University (WVU) and started enrolling students in fall 2018. Soon after establishing these programs, the standalone project Attracting and Cultivating Cybersecurity Experts and Scholars through Scholarships (ACCESS) was funded by the NSF S-STEM program. ACCESS objectives are: (1) increase the annual enrollment in the B.S. and AoE in Cybersecurity; (2) provide co-curricular activities intended to bolster students' academic achievement and career prospects; (3) establish partnerships with cybersecurity employers from the private and public sector; and (4) explore the impact of the ACCESS activities on students' academic success and professional careers.

Over six years, a total of 227 semester scholarships were awarded to 63 unique students from five cohorts. The successful recruitment resulted from a wide range of outreach activities tailored to different student identities, across different academic stages. 35% of new scholars were recruited while they were still in high school. ACCESS has contributed to the increased enrollment in the B.S. and AoE in Cybersecurity at WVU - from 50 students (US citizens) at the start of the project in spring 2020 to 201 students (US citizens) in spring 2025. To date, out of the 63 ACCESS scholars, 31 scholars have graduated and additional 25 scholars are still pursuing their degree.

The ACCESS project goes beyond awarding scholarships. To aid students' success, the project team developed and offered numerous co-curricular activities and support services. These include social events which enabled community building. Furthermore, all scholars received faculty mentoring, most scholars participated in the CyberWVU student organization which helped them develop hands-on skills, and some scholars were involved in undergraduate research.

ACCESS has strengthened the existing and created new partnerships with many cybersecurity employers. The ACCESS team has organized 26 seminars that were given by renowned cybersecurity experts. These events allowed students to gain practical knowledge, improve their soft skills, develop professional networks, and learn about different career paths and opportunities.

Lastly, the research team investigated the impact of the ACCESS activities on students' success. For instance, an anonymous survey showed that all scholars strongly or somewhat agreed that they learned about career opportunities they would not have otherwise known about, were more confident about starting a career in cybersecurity, had the opportunity to learn from cybersecurity professionals, and were able to access resources that will help them in their field.

Keywords: Cybersecurity education, scholarships, co-curricular activities, mentoring, seminars, educational research.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

The OverClock Experience HACKnet: Living and learning in the modern cybersecurity residence hall

[Poster]

David Richards, Grand Canyon University, AZ, david.richards1@gcu.edu

Extended Abstract

This session examines Grand Canyon University's "OverClock Experience" initiative. An innovative program designed to advance cybersecurity education with an immersive, continuous, and gamified network environment embedded directly within campus residential housing. The initiative is distinguished by its 24/7 access for students living in a residence hall and the integration of experiential learning methodologies with simulated threat landscapes and cross-disciplinary collaboration, thereby cultivating a robust ecosystem for sustained skill development. A prominent feature of The OverClock Experience HACKnet is its tiered network design, which allows participants to engage in authentic cybersecurity activities, including network scanning, penetration testing, and offensive security operations that are typically restricted in operational environments.

Network challenges are structured on three progressive levels: "Easy" networks provide hackable web applications and virtual machines; "Medium" networks mimic corporate infrastructures with lax security controls and extended patching intervals; and "Hard" networks simulate enterprise networks with stringent security controls and frequent patching schedules. Points are assigned based on the complexity of security breaches encountered—entry-level attacks garner basic points, while more sophisticated exploits yield greater rewards. Students are tasked with identifying and submitting digital flags as evidence of successful breaches, fostering systematic analysis and strategic thinking.

All network environments are equipped with industry-standard security monitoring tools, affording students valuable experience in security operations and incident response. This hands-on access within the residential setting enables learners to reinforce theoretical instruction with immediate practical application, effectively bridging the gap between classroom learning and the operational realities of modern cybersecurity. Every participant must sign a rigorous AUP, defining the clear boundaries between sanctioned research and prohibited activity. All HACKnet activities are governed by the University's overarching code of conduct, ensuring that technical exploration remains rooted in community values. Attendees will gain insights into the implementation and pedagogical merits of continuous, experiential learning, and appreciate how the "living and learning" model exemplified by The OverClock HACKnet prepares students for the complex, evolving demands of the cybersecurity field.

Keywords: Student activities, ethical sandbox, competitions, workforce readiness, cybersecurity.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

AI based dynamic threat modeling for assessing access control posture in cyber physical systems

[Poster]

Indrakshi Ray, Colorado State University, CO, indrakshi.ray@colostate.edu

Shwetha Gowdanakatte, Colorado State University, CO, shwetha.gowdanakatte@colostate.edu

Extended Abstract

Cyber-Physical Systems (CPS) and Industrial Control System (ICS) environments now depend on advanced access-control mechanisms to manage digital and physical interactions. With the adoption of technologies like IEC 61850 automation and AI integration, new threats emerge from RBAC/ABAC misconfigurations and risky cross-zone policies—risks that traditional threat modeling fails to capture. To address this, we introduce a dynamic, access-control-centric threat-modeling framework that combines automated policy analysis, hybrid STRIDE-ICS mapping, attack graphs, AI planning, and simulation.

The approach starts by modeling CPS architecture and identifying access-control vulnerabilities, mapping them to MITRE ATT&CK for ICS tactics. Further analysis uses STRIDE and ICS semantics to formalize each vulnerability into preconditions, actions, and outcomes. These elements are then used to generate attack graphs, which are transformed into PDDL-based AI planning models for generating optimized attack scenarios. A simulation testbed validates these scenarios in a realistic digital substation environment.

A continuous feedback loop refines models with results from simulations, adapting to new attack strategies and operational changes. The final phase assesses performance based on attack success, plan quality, detection coverage, and comparison to traditional models. This work presents an AI-enabled, access-control-focused, dynamic threat-modeling framework that unifies ATT&CK semantics, automated planning, and simulation to enhance proactive defense for modern OT systems.

Novel contributions include: (1) elevating authorization as a primary attack surface for OT, (2) action-level mapping to ATT&CK ICS to ground findings in standardized TTPs, and (3) AI-planning readiness by emitting initial states and actions directly from policies—reducing hand-crafted modeling burden in later phases. This positions defenders to prioritize mitigations with evidence that is both standards-aligned and operationally verifiable within a realistic digital substation context.

Keywords: Digital substation, IEC 61850 (MMS/GOOSE), access control, ABAC/RBAC, authorization graph, MITRE ATT&CK for ICS, threat modeling, AI planning, attack graphs.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

SECURING AI systems against data, model, and tool poisoning attacks

[Poster]

Meryem Abouali, John Jay College of Criminal Justice, NY, mabouali@jjay.cuny.edy

Danish Merchant, John Jay College of Criminal Justice, NY, mohammeddanish.merchant63@login.cuny.edu

Mauricio Embus Perez, John Jay College of Criminal Justice, NY, Mauricio.embusperez24@login.cuny.edu

Perla Dahan, John Jay College of Criminal Justice, NY, perla.dahan07@login.cuny.edu

Extended Abstract

The rapid deployment of large language models (LLMs) and AI-driven agents across cybersecurity, education, and critical infrastructure has introduced new threats that are not adequately addressed by traditional security models. Among the most pressing risks are data poisoning, model manipulation, prompt injection, and tool-level exploitation, which undermine the integrity and trustworthiness of AI systems. These challenges are recognized in the OWASP LLM Top 10, particularly under LLM04: Data and Model Poisoning. This extended abstract presents an empirical study of AI poisoning attacks and defense mechanisms relevant to the CAE-CD community's focus on applied cyber defense.

This project evaluates AI poisoning across three stages of the AI lifecycle: training data, runtime inference, and AI agent tool integration. At the data level, multiple poisoning techniques, including backdoor insertion, label manipulation, bias injection, and toxic content poisoning, were applied to training datasets to examine their impact on model behavior. The results demonstrate that compromised data can significantly distort outputs while remaining difficult to detect through conventional validation methods.

At the inference stage, the study examines prompt injection, context poisoning, and jailbreak attacks against a deployed LLM environment, showing that models remain vulnerable after deployment when exposed to adversarial inputs. The project also investigates Model Context Protocol (MCP) tool poisoning, an emerging attack surface in AI agent ecosystems. Manipulated tool responses and contextual inputs were shown to influence agent behavior without modifying the underlying model.

To mitigate these risks, a layered defense framework was evaluated, combining dataset validation, anomaly detection, prompt filtering, runtime monitoring, and policy-based controls. The findings demonstrate that defense-in-depth strategies substantially improve resilience to both training-time and runtime attacks and highlight the need to integrate AI-specific threat models into cyber defense education and practice.

Keywords: AI security, data poisoning, large language models, OWASP LLM04, cyber defense education, prompt injection.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Strengthening home networks: Evaluating routers against NIST standards

[Poster]

Oluwapemiisin G. Akingbola, Morgan State University, MD, olaki62@morgan.edu

Toluwalogo B. Odumosu, Morgan State University, MD, toluwalogo.odumosu@morgan.edu

Extended Abstract

The state of a home router plays a crucial role in determining how secure a home network connection is. As the primary gateway through which all network traffic passes, routers represent a significant point of control and a prime target for attackers seeking to gain unauthorized access to connected devices and individuals' personal data. With the increasing number of connected smart devices and other IoT devices in the household, modern routers need to prioritize safety.

Despite the significant increase in features and capabilities in modern home routers (e.g., mesh networking, gaming optimization), there is limited research that evaluates whether these routers comply with proposed or established cybersecurity policies. This study addresses this gap by focusing on the NIST IR 8425 and NIST IR 8425A standards. Fundamental requirements such as access control, authentication, firmware practices, and network segmentation capabilities are reviewed for each router placed under consideration in this study.

The results show significant variation in home routers' compliance with the proposed framework, highlighting how much work the market still needs to do to meet these proposed minimum standards. This gap highlights the need for more robust policy oversight to improve overall home network security practices.

Keywords: Home network security, consumer-grade routers, router security evaluation, NIST cybersecurity standards, NIST IR 8425, NIST IR 8425A, Internet of Things security, IoT-integrated home networks.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Defense (CD) Track

A Publication of the CAE in Cybersecurity Community National Center

Gated cross-attention matcher for aligning courses with relevant KUs

[Poster]

Aashish Pandey, University of North Texas, TX, aashishpandey2@my.unt.edu

Ram Dantu, University of North Texas, TX, ram.dantu@unt.edu

Extended Abstract

Curriculum alignment remains a persistent challenge in academic program design, accreditation processes, and assessment development. Course syllabi and outcome tables typically describe instructional intent at a broad level, whereas examinations evaluate fine-grained competencies through micro-assessments such as short-answer prompts, sub-topics, and specific skill checks. This cross-granularity mismatch, combined with inconsistent phrasing and boilerplate educational language, makes traditional keyword search and generic embedding retrieval methods unreliable.

To address this problem, we propose a retrieval-based alignment framework centered on a lightweight Gated Cross-Attention Module (GCAM) that improves robustness, adaptability, and interpretability. The system begins with a structured data layer that ingests raw curriculum artifacts, including syllabi and learning outcome tables. These documents are parsed and normalized into Knowledge Units, each represented by a title, a list of topics, and associated learning outcomes. The model layer employs a frozen large language model encoder to produce contextual token representations for both queries and memory items. Rather than fine-tuning the full backbone, GCAM introduces lightweight adaptation through low-rank projection layers and two interpretable gating mechanisms.

Cross-attention is applied between query tokens and memory tokens to generate a query-aware memory representation. The Token Gate selectively suppresses irrelevant or noisy memory tokens while emphasizing concept-bearing terms. The Head Gate dynamically activates the most useful attention heads for a given query, allowing the model to adapt its alignment behavior depending on query intent and granularity. Together, these mechanisms convert standard cross-attention into an intent-adaptive and noise-resistant semantic matching module.

The model is trained using a contrastive objective that encourages aligned pairs to produce stronger similarity scores than non-aligned pairs. During inference, course titles, topics, or micro-assessments are scored against the KU memory bank to retrieve the most relevant Knowledge Units and associated learning outcomes. This retrieval-based framework reduces manual alignment effort, improves consistency across instructors, and enhances interpretability through explicit gating behavior.

Keywords: Curriculum alignment, knowledge units, cross-attention, contrastive learning, educational retrieval, semantic matching, gating mechanism.

CAE-R RESEARCH

R Track: Program Committee Co-Chairs

Tommy Morris
The University of Alabama
in Huntsville, AL
tommy.morris@uah.edu

Isabelle Brown-Cantrell
The University of Alabama
in Huntsville, AL
isabelle.brown@uah.edu

Mohammad Einaam Alim
The University of Alabama
in Huntsville, AL
mohammad.alim@uah.edu

R Track: Proceedings Editorial Preface

2026 CAE in Cybersecurity Community Symposium
National Centers of Academic Excellence in Cybersecurity (NCAE-C)
Pittsburgh, PA

CAE-R-designated institutions play a distinctive role in advancing the scientific foundations of cybersecurity through original research, mentorship of emerging scholars, and cultivation of faculty expertise that pushes the boundaries of what is known and what is possible in defending our national infrastructure. The CAE-R institutions serve as powerhouses of discovery with their faculty and students conducting research that spans areas such as network security, cryptography, formal methods, artificial intelligence for cyber defense, privacy-preserving computation, and cyber-physical systems security, among many others.

The 2026 CAE Symposium Research Track investigates the relationship between artificial intelligence and cybersecurity, examining how AI and machine learning techniques can transform threat detection, vulnerability assessment, and autonomous defensive operations, while confronting the equally critical challenge of hardening these intelligent systems against exploitation. The track extends well beyond pure research, placing substantial emphasis on cultivating the national cybersecurity talent pipeline through novel AI- and ML-centric pedagogical approaches. Accepted presentations cover curriculum innovation, immersive training platforms, and actionable strategies for closing the faculty expertise gap in these rapidly evolving disciplines. In parallel, the program hosts dedicated "Getting to Know a CAE" and "Getting to Know a PhD" sessions that highlight institutional research achievements and the dissertation contributions of emerging scholars within the CAE-R community. This year marked another milestone for the NCAE-C Community. From a competitive pool of 31 submissions, the program committee curated a program of 10 accepted papers, two institutional highlight talks, and one outstanding PhD scholar presentation. Taken together, these elements create a collaborative venue for disseminating proven practices, exchanging hard-won lessons, and charting future directions for the security challenges that span the entire CAE network.

The vision for this year's community symposium was to organize an academic double-blind peer-review process using the EasyChair platform for presentation submissions and produce a proceedings book to document all the hard work by community members. We strongly believe in the value of coming together to share best practices, stories of successful collaborations, lessons learned, future directions, and solutions to community-wide issues.

We would also like to thank Dr. Tony Coulson, Amy Hysell, and their team at the CAE in Cybersecurity Community National Center (https://caecommunity.org/) for supporting the CAE Community of Practice in Cyber Research (CoP-R) (https://www.caecommunity.org/cop-cyber-research) and the 2026 CAE in Cybersecurity Community Symposium.

The 2026 CAE Community Symposium - R Track Program Committee Co-Chairs,

Tommy Morris, Ph.D. ***Isabelle Brown-Cantrell*** ***Mohammad Einaam Alim***
University of Alabama in Huntsville, AL

R Track: Refereed Extended Abstract Proceedings for Presentations

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Research (R) Track

A Publication of the CAE in Cybersecurity Community National Center

RAG-targeted Adversarial Attack on LLM-based threat Detection and Mitigation Framework in IoT

[Presentation]

Maanak Gupta, Tennessee Tech University, TN, mgupta@tntech.edu

Seif Ikbarieh, Tennessee Tech University, TN, sikbarieh@tntech.edu

Kshitiz Aryal, University of Nebraska Omaha, NE, karyal@omaha.edu

Extended Abstract

The rapid expansion of the Internet of Things (IoT) is reshaping communication and operational practices across industries, but it also broadens the attack surface and increases susceptibility to security breaches. Artificial Intelligence has become a valuable solution in securing IoT networks, with Large Language Models (LLMs) enabling automated attack behavior analysis and mitigation suggestion in Network Intrusion Detection Systems (NIDS). Despite advancements, the use of LLMs in such systems further expands the attack surface, putting entire networks at risk by introducing vulnerabilities such as prompt injection and data poisoning.

In this work, we attack an LLM-based IoT attack analysis and mitigation framework to test its adversarial robustness. We construct an attack description dataset and use it in a targeted data poisoning attack that applies word-level, meaning-preserving perturbations to corrupt the Retrieval-Augmented Generation (RAG) knowledge base of the framework. We then compare pre-attack and post-attack mitigation responses from the target model, ChatGPT-5 Thinking, to measure the impact of the attack on model performance, using an established evaluation rubric designed for human experts and judge LLMs. Our results show that small perturbations degrade LLM performance by weakening the linkage between observed network traffic features and attack behavior, and by reducing the specificity and practicality of recommended mitigations for resource-constrained devices.

Keywords: Adversarial Poisoning Attack, Internet of Things (IoT) Security, Large Language Model, Retrieval-Augmented Generation, Cybersecurity.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Research (R) Track

A Publication of the CAE in Cybersecurity Community National Center

A multi-agent system for enhancing static taint analysis of JavaScript applications

[Presentation]

William Enck, North Carolina State University, NC, whenck@ncsu.edu

Extended Abstract

JavaScript is one of the most widely used programming languages, with the npm registry containing over 3 million JavaScript packages. Unfortunately, existing Static Application Security Testing (SAST) tools fail to effectively discover vulnerabilities in JavaScript due to (1) dynamic features that complicate data flow extraction and (2) npm's large library ecosystem that makes it difficult to identify relevant sources/sinks and establish taint propagation across dependencies.

In this presentation, we propose a multi-agent system that strategically combines the semantic understanding of Large Language Models (LLMs) with traditional static program analysis to extract taint specifications, including sources, sinks, call edges, and library flow summaries tailored to each package. Conceptually, we use static program analysis to calculate a call graph and defer to an LLM to resolve call edges that cannot be resolved statically, as well as identify per-package taint sources and sinks. We use three specialized LLM agents, which enhance vulnerability analysis in complementary ways. First, the *Source/Sink Agent* identifies potential sources and sinks for vulnerability categories guided by the corresponding MITRE CWE description. Second, the *CallGraph Agent* iteratively resolves call edges that the SAST tool cannot determine statically, producing both first-party edges that connect calls to callees within the package and candidate flow summaries for third-party dependencies that conservatively assume taint propagation. Third, the *Flow Summary Agent* refines any candidate flow summaries present in vulnerability paths, determining whether the corresponding APIs sanitize or propagate taint, filtering false positives while preserving true vulnerabilities. The resulting taint specification (sources, sinks, and call edges) is provided to the SAST tool to complete the vulnerability analysis.

We integrated our approach with CodeQL, a state-of-the-art SAST tool, and demonstrated its effectiveness by detecting 109 of 162 vulnerabilities previously undetectable by CodeQL. Furthermore, we find 7 novel vulnerabilities spanning 6 popular npm packages. In doing so, we demonstrate that LLMs can practically enhance existing static program analysis algorithms, combining the strengths of both symbolic reasoning and semantic understanding for improved vulnerability detection.

Keywords: Static taint analysis, vulnerability detection, Javascript, LLMs, multi-agent systems.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Research (R) Track

A Publication of the CAE in Cybersecurity Community National Center

Steganography with large language models: Key sensitivity analysis

[Presentation]

Alexander V. Mantzaris, University of Central Florida, FL, Alexander.Mantzaris@ucf.edu

Wissam Ghantous, University of Central Florida, FL, wissam.ghantous@ucf.edu

Extended Abstract

Steganographic methods based on large language models (LLMs) enable the embedding of hidden information within fluent, human-readable text, typically controlled by a secret prompt or seed acting as a key. Rank-based LLM constructions have recently emerged as particularly effective, offering both high embedding rates and strong distributional indistinguishability. However, the extent to which their outputs depend on the precise choice of key has not been systematically characterized.

We investigate this question through a series of concrete case studies drawn from a representative rank-based LLM stegosystem inspired by the work of Norelli and Bronstein. Our analysis centers on how small perturbations of the key influence the generated text, using multiple notions of textual separation, including token-wise disagreement patterns and aggregate distance measures. Experiments are conducted with a fixed language model, combining synthetic control prompts with passages from standard public-domain texts. Across all examined scenarios, we observe that minor changes to the key rapidly decorrelate the resulting stegotexts, driving measured distances close to their maximal values and exhibiting little sensitivity to the magnitude of the perturbation. These findings indicate that the key-to-output map behaves analogously to a cryptographic primitive with respect to key variation, providing empirical evidence for robustness against attacks that exploit key proximity and emphasizing the importance of exact key specification.

Keywords: Steganography, large language models, rank-based encoding, key sensitivity.

References:

Norelli, A., & Bronstein, M. (2025). LLMs can hide text in other text of the same length. *arXiv preprint arXiv:2510.20075*.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Research (R) Track

A Publication of the CAE in Cybersecurity Community National Center

Integrating GenAI and the DCWF into a graduate cybersecurity course: A framework for prompt-based auditing and risk mitigation

[Presentation]

Yair Levy, Nova Southeastern University, FL, levyy@nova.edu

Extended Abstract

As Generative Artificial Intelligence (GenAI) continues to reshape the technological landscape, cybersecurity professionals must adapt to leverage these tools for cybersecurity including risk management, auditing, and network hardening. This presentation will outline a pedagogical framework for integrating GenAI into a graduate-level cybersecurity course, based on a micromodule recently published on the CAE's CLARK center platform. The assignment challenges students to bridge the gap between theoretical frameworks—specifically the United States (U.S.) Department of Defense (DoD) Cyber Workforce Framework (DCWF) (2024)—and the practical application of AI-driven automation.

The assignment requires students to identify specific DCWF work roles and tasks, such as Cyber Defense Analyst (Code 511), and develop advanced prompt engineering skills to generate realistic operational scenarios. A critical component of this assignment is the multi-platform auditing phase, where students use at least three distinct GenAI tools to cross-evaluate the initial output for relevancy, accuracy, and compliance with national standards like NIST Cybersecurity Framework 2.0 or CMMC 2.0. This approach is grounded in recent research demonstrating the efficacy of comparing GenAI platforms to mitigate cybersecurity risks (Gafni & Levy, 2025). By adopting this assignment, faculty members can provide students with hands-on experience in prompt-based auditing and scenario evaluation, essential skills for modern cybersecurity risk management. The talk will focus on the practicalities of implementing this graduate micromodule assignment in a course setting, including guidelines for maintaining academic integrity and professional formatting standards. Attendees will leave with a proven roadmap for preparing students to use AI for cybersecurity and not just as a productivity tool, but as a robust mechanism for increasing the cybersecurity posture of an organization with the ability to link the results of national standards.

Keywords: Generative AI (GenAI) for Cybersecurity, DCWF in graduate academic assignments, Cyber risk mitigation, Prompt development for risk assessment.

References:

Gafni, R., & Levy, Y. (2025). Comparing GenAI platforms on cybersecurity management task performances. *Information and Computer Security, 33*(4). https://doi.org/10.1108/ICS-05-2024-0130

U.S. Department of Defense. (2024). *DoD cyber workforce framework (DCWF).* https://www.cyber.mil/dod-workforce-innovation-directorate/dod-cyber-workforce-framework/dcwf

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Research (R) Track

A Publication of the CAE in Cybersecurity Community National Center

Assessing and ensuring green traffic realism in cyber ranges and competitions

[Presentation]

Steven Lamp, University of Virginia, VA, vxn3k@virginia.edu

Jason D. Hiser, University of Virginia, VA, jdh8d@virginia.edu

Anh Nguyen-Tuong, University of Virginia, VA, an7s@virginia.edu

Jack W. Davidson, University of Virginia, VA, jwd@virginia.edu

Extended Abstract

Cyber ranges and collegiate cyber competitions depend on realistic network environments to support effective training, evaluation, and research. While red team (attacker) and blue team (defender) activity receive significant attention, the realism of *green traffic*—background activity generated by benign users, services, and systems—is often assumed rather than validated. In practice, green traffic can strongly influence defender perception, alert fidelity, and exercise outcomes. Unrealistic background behavior can distort defensive strategies, obscure meaningful attack signals, and reduce the training value of environments used for competitions such as NCCDC, MACCDC, and CPTC. Despite its importance, there is currently no widely adopted method for measuring whether green traffic reflects human-like behavior or for validating whether attempted improvements are effective across different operational contexts.

We present PHASE (Passive Human Activity Simulation Evaluation), a passive, measurement-based framework for assessing and ensuring the realism of green traffic in cyber ranges and competition environments. PHASE uses time-series machine learning models trained on real-world network data to identify temporal and behavioral patterns characteristic of human activity. Across eight diverse datasets—including network data from large operational networks and publicly available competition datasets—PHASE achieves an average classification accuracy of approximately 91%, with strong balanced accuracy and precision exceeding 90%, demonstrating reliable identification of human-predominant behavior while limiting false positives. PHASE operates entirely on passively collected network connection logs and scores activity along a spectrum from human-predominant to automation-predominant behavior. Because PHASE relies solely on existing network telemetry, it introduces no probes or artifacts that could influence observed behavior, making it suitable for both operational ranges and competitive settings.

In addition to producing quantitative scores, PHASE incorporates explainability techniques that reveal *why* observed traffic is classified as human-like or non-human-like. Using these insights, PHASE was applied to evaluate a widely used green traffic generation tool. A simple configuration change—introducing realistic idle periods—improved measured behavioral realism by up to 55 percentage points. By enabling evidence-based evaluation of background traffic, PHASE supports the development of more authentic cyber ranges and competitions that better reflect real-world operational conditions, thereby improving cybersecurity training outcomes.

Keywords: Network traffic analysis, Deep neural network (DNN), Green traffic, Cyber range.

Understanding and defending against data and model poisoning

[Presentation]

Qian Lou, University of Central Florida, FL, qian.lou@ucf.edu

Extended Abstract

Modern AI systems are increasingly deployed in security- and safety-critical settings, yet they remain vulnerable to data and model poisoning attacks that can silently compromise reliability, fairness, and trust. This talk presents a unified view of my research on understanding and defending against such threats across the AI lifecycle, combining red-team–driven risk identification with principled, information-theoretic defenses.

I will describe a series of benchmark studies and threat models that expose previously unrecognized vulnerabilities in contemporary AI systems. These include inference-time backdoor attacks in text and vision models (Lou et al., 2023), training-time data poisoning, and the first systematic Trojan and prompt-injection attacks on large language models (Xue et al., 2023). Beyond conventional model pipelines, I will discuss emerging risks in reasoning frameworks, retrieval-augmented generation, and even hardware-level bit-manipulation attacks, illustrating how integrity failures can propagate across the entire AI stack. On the defense side, I will present mitigation frameworks that move beyond heuristic detection toward proactive and provable robustness. This includes self-supervised techniques (Zheng et al., 2024) that detect and remove Trojan triggers without labeled data, as well as information-theoretic defenses that provide certified guarantees (Lou et al., 2024) against universal and adversarial perturbations in large language models. Together, these approaches establish a principled foundation for building AI systems that are secure, resilient, and trustworthy in real-world deployments.

Keywords: AI Security, Data and Model Poisoning, Large Language Models, Provable Robustness

References:

Lou, Q., Liu, Y., & Feng, B. (2023). TrojText: Test-time invisible textual trojan insertion. *Proceedings of the Eleventh International Conference on Learning Representations*.

Lou, Q., Liang, X., Xue, J., Zhang, Y., Xie, R., & Zheng, M. (2024). CR-UTP: Certified robustness against universal text perturbations on large language models. *In Findings of the Association for Computational Linguistics: ACL 2024* (pp. 9863-9875).

Xue, J., Zheng, M., Hua, T., Shen, Y., Liu, Y., Bölöni, L., & Lou, Q. (2023). Trojllm: A black-box trojan prompt attack on large language models. *Advances in Neural Information Processing Systems, 36*, 65665-65677.

Zheng, M., Xue, J., Wang, Z., Chen, X., Lou, Q., Jiang, L., & Wang, X. (2024). Ssl-cleanse: Trojan detection and mitigation in self-supervised learning. *Proceedings of the European Conference on Computer Vision* (pp. 405-421).

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Research (R) Track

A Publication of the CAE in Cybersecurity Community National Center

PhishGauge: Visual phishing detection with generative augmentation

[Presentation]

Tam V. Nguyen, University of Dayton, OH, tamnguyen@udayton.edu

Zhongmei Yao, University of Dayton, OH, zyao01@udayton.edu

Ju Shen, University of Dayton, OH, jshen1@udayton.edu

Extended Abstract

Phishing websites pose a significant cybersecurity threat by visually imitating legitimate online services to deceive users into revealing sensitive information. Most existing phishing detection systems rely textual features such as HTML, JavaScript source code or URL (Popescul & Radu, 2025). However, such information is not always accessible due to encryption, obfuscation, or client-side deployment constraints. In this work, we introduce PhishGauge, short form of **Phish**ing Detection with **G**enerative **Aug**m**e**ntation, that relies solely on visual features extracted from website screenshots and leverages Generative Artificial Intelligence (GenAI) to improve detection performance and robustness. Our proposed PhishGauge framework formulates phishing detection as an image-based classification problem using rendered webpage screenshots. Here, the model is trained to learn discriminative visual patterns observed in phishing websites. To address the limited availability and rapid evolution of phishing website datasets, we adopt GenAI models (Rombach et al., 2022) to synthesize realistic webpage screenshots for data augmentation. By enriching training data with diverse, visually plausible phishing examples, our work improves the performance of visual phishing website detection via the experiments on Phish-IRIS dataset (Dalgic et al., 2018).

Keywords: Visual phishing detection, generative AI, data augmentation, cybersecurity.

References:

Dalgic, F. C., Bozkir, A. S., & Aydos, M. (2018). Phish-iris: A new approach for vision based brand prediction of phishing web pages via compact visual descriptors. *Proceedings of the 2018 International Symposium on Multidisciplinary Studies and Innovative Technologies (ISMSIT)* (pp. 1-8).

Popescul, D., & Radu, L. D. (2025). AI in phishing detection: a bibliometric review. *Frontiers in Artificial Intelligence*, *8*, 1496580.

Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2022). High-resolution image synthesis with latent diffusion models. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 10684-10695).

Zhang, L., Rao, A., & Agrawala, M. (2023). Adding conditional control to text-to-image diffusion models. *Proceedings of the IEEE/CVF International Conference on Computer Vision* (pp. 3836-3847).

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Research (R) Track

A Publication of the CAE in Cybersecurity Community National Center

Deliberative reasoning with system-2 agents for deep-logic vulnerability discovery and exploitation

[Presentation]

Xiuwen Liu, Florida State University, FL, xliu@fsu.edu

Extended Abstract

The success of AI built on Large Language Models (LLMs) has enabled highly automated agents for penetration testing and Capture-the-Flag (CTF) challenges. For instance, agents can now automatically plan and execute exploits on targets like Metasploitable given only an IP address, or solve common CTF problems by simply accessing relevant file directories and standard tools. This success pushes the frontier toward developing novel complex exploits and discovering zero-day vulnerabilities.

However, scaling LLMs to find deep-logic vulnerabilities presents significant technical hurdles. While "System-1" thinking—relying on pattern matching and memorization—suffices for benchmarks with ample testing samples, structural limitations render this paradigm inherently deficient for complex problems requiring extensive reasoning. For example, LLMs rely on superpositions to represent concepts efficiently; this introduces cross-talk and interference patterns that manifest as hallucinations and persistent errors. When chaining many steps, overall performance degrades rapidly, often reaching zero after just two hundred steps. While error rates may improve, these structural issues make removing errors practically impossible.

Specific to program analysis, dependency structures in vulnerabilities differ vastly from natural language, yet optimal representations and dependency models remain missing. Additionally, chaining the hundreds or even millions of steps required to identify deep-logic vulnerabilities remains an open challenge. Furthermore, current LLMs lack the "lifelong learning" capabilities of human experts, as they stop learning after the training phase.

To address these limitations, we are developing a novel agentic AI system. **1)** Recognizing that knowledge in existing writeups is poorly utilized by next-token prediction, we employ a sample-driven top-down synthesis to discover parameterized patterns defined by the chaining of atomic primitives. **2)** We design a structured "System-2" architecture integrating Monte-Carlo tree search to enable long-horizon reasoning. **3)** We implement an agentic memory system to accumulate experience and enable lifelong learning. Together, this system establishes a new frontier where zero-day vulnerabilities and exploits are discovered efficiently.

Keywords: Agentic AI, Large Language Models, System-2 Reasoning, Vulnerability Discovery, Automated Exploit Generation, Program Synthesis.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Research (R) Track

A Publication of the CAE in Cybersecurity Community National Center

Artificial intelligence driven digital twin and adaptive autonomy for safe and secure UAV operations

[Presentation]

Md Tauhidur Rahman, Florida International University, FL, mdtrahma@fiu.edu

Shafika Showkat Moni, Embry Ridle Aeronautical University, FL, monis@erau.edu

M. Ilhan Akbas, Embry Ridle Aeronautical University, FL, akbasm@erau.edu

Extended Abstract

Conventional Unmanned Aerial Vehicle (UAV) planning and evaluation strategies handle flight safety and cybersecurity independently, despite their intrinsic interdependence. In real-world operations, particularly in contested or GPS-denied environments, a cyberattack can directly impair onboard safety functions like obstacle avoidance, while loss of situational awareness may increase vulnerability to further attacks. Furthermore, hardware-based attacks can compromise critical sensing or control components, resulting in degraded performance, loss of flight stability, or complete mission failure, thereby exacerbating both safety risks and cybersecurity vulnerabilities in tightly coupled UAV operations. Moreover, most UAVs operate as part of a fleet or swarm, where cascading failures across multiple agents can compromise mission objectives and coordination.

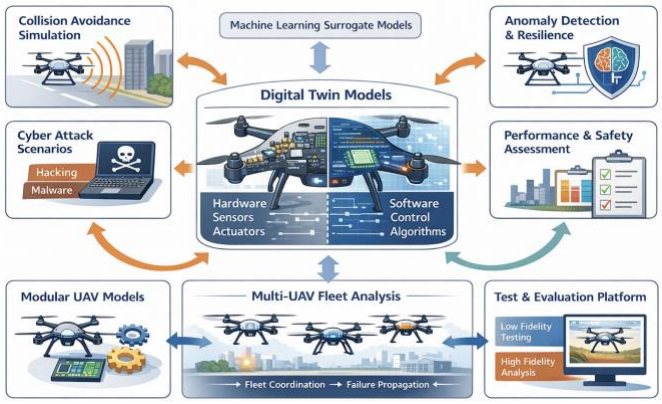

Fig. 1: Adaptive Autonomy for Safe and Secure UAV Operations using AI-driven Digital Twin.

To address these challenges, this research proposes, as shown in Fig. 1, an Artificial Intelligence (AI) Driven Digital Twin and Adaptive Autonomy approach that blends digital engineering, real-time autonomy adaptation, and adversarial scenario modeling. Our proposed system supports scalable modular simulation of UAVs, each paired with a high-fidelity digital twin capable of representing UAV hardware, software, sensor payloads and flight dynamics. The system will facilitate 1) Assessment of hardware/software vulnerabilities via synthetic cyber-physical attacks, 2) Planning and maneuvering strategies for UAVs, 3) Analysis of multi agent coordination under spoofing and degraded GPS or communication, 4) Real-time detection of anomalies and autonomous fallback control mechanisms.

Keywords: AI-driven Digital Twin, Digital Twin of UAV, Safety and Security of UAV operations.

R Track: Refereed Extended Abstract Proceedings for INSuRE Presentations

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Research (R) Track

A Publication of the CAE in Cybersecurity Community National Center

A machine-learning based approach to malicious document detection for RAG chunk ingestion

[INSuRE Presentation]

Bryan Chan, Stevens Institute of Technology, NJ, bchan4@stevens.edu

Extended Abstract

Retrieval Augmented Generation (RAG) is used commonly in modern enterprise systems for large language models (LLMs), but this new technology comes with security hazards. One such exploit that is created from these hazards is malicious document injection - the ability to compromise an information base that ends up feeding misinformation to an LLM. To prevent this exploit from happening, we suggest a baseline naive classifier that can detect and prevent these malicious injections.

Keywords: Malware detection, LLM, RAG, injection.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Research (R) Track

A Publication of the CAE in Cybersecurity Community National Center

WinMango: Extending automated static binary analysis to windows PE binaries

[INSuRE Presentation]

Sam Mabry, University of Alabama at Birmingham, AL, smabry@uab.edu

Extended Abstract

This presentation covers the development and evaluation of WinMango, an extension of Operation Mango (USENIX Security 2024) that enables automated vulnerability detection in Windows PE binaries using WinAPI-specific taint analysis. We discuss the engineering challenges of adapting a Linux ELF-focused framework to handle PE/IAT indirect calls, MSVC and MinGW compiler support, and PDB symbol resolution. We present evaluation results across 4,700+ synthetic test binaries derived from the NIST Juliet Test Suite, covering 16 WinAPI source functions, 11 sink functions, and 13 data-flow variants across both compilers. We also discuss remaining limitations and the path toward real-world malware analysis.

Keywords: Static binary analysis, malware detection.

R Track: Refereed Extended Abstract Proceedings for PhD Student Highlight Talk

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Research (R) Track

A Publication of the CAE in Cybersecurity Community National Center

Efficiently finding aliasing bugs in multilanguage Rust applications

[PhD Student Highlight Talk]

Hanan Hibshi, Carnegie Mellon University, PA, hhibshi@cmu.edu

Ian McCormack, Carnegie Mellon University, PA, icmccorm@andrew.cmu

Extended Abstract

Most critical security vulnerabilities are caused by memory safety issues. To counter this threat, developers are increasingly transitioning to writing new, security-critical software in Rust: a programming language that provides inherent safety guarantees. However, new Rust components often need to interoperate with legacy systems written in C and C++. These languages lack the restrictions that Rust needs to prevent safety issues, making it possible to trigger security vulnerabilities in otherwise *safe* components via foreign function calls.

Foreign functions are one of several "unsafe" features that are allowed to bypass Rust's static restrictions on aliasing and mutability. Developers who use these features still need to ensure that they are following the rules of Rust's evolving aliasing model. Otherwise, their programs can behave erratically, reintroducing the types of security issues that Rust was designed to prevent. Miri, a popular bug-finding tool, is the only method that developers can use to find aliasing violations. However, Miri has high overhead and limited support for foreign function calls. Meanwhile, Rust is being adopted in security-critical C++ applications like Chromium, Android, and the Linux Kernel, which require interoperation with C and C++. Miri is not nearly as useful for these applications, leaving developers without a solution for finding Rust-specific bugs in these contexts.

We are creating BorrowSanitizer: a new dynamic analysis tool for finding Rust-specific aliasing bugs in these kinds of multilanguage applications. Unlike Miri, our approach relies on inserting run-time checks ahead-of-time, during compilation. We extract Rust-specific aliasing information and lower it into the Rust compiler's LLVM backend, which is shared with C and C++. At this level, we insert instrumentation that tracks the "provenance" metadata associated with pointers. Our run-time assertions validate this metadata prior to each memory access.
In preliminary research, we created a prototype extension to Miri and demonstrated that multilanguage bugs exist in high-profile Rust libraries. We presented these results at the 2025 International Conference on Software Engineering. Now, with BorrowSanitizer, we are building an entirely new tool from the ground up. Throughout 2026, we will be transitioning from an initial research prototype into a production-ready solution. Our ongoing development has been supported by funding from Google and the Rust Foundation, and we are collaborating with the Rust Team through an ongoing Project Goal.

Keywords: Rust, secure programming, bug finding, software security.

R Track: Refereed Extended Abstract Proceedings for School Highlight Talks

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Research (R) Track

A Publication of the CAE in Cybersecurity Community National Center

Get to know a CAE: Washington State University cybersecurity program: Developing the next-generation cyber workforce

[School Highlight Talk]

Assefaw Gebremedhin, Washington State University, WA,
assefaw.gebremedhin@wsu.edu

Extended Abstract

Washington State University (WSU) was designated as CAE-R school in 2025. A few strengths distinguish WSU as a key player at the forefront of modern cybersecurity education. With support from the Washington state government, a new Bachelor of Science in Cybersecurity with emphasis on cyber operations has been launched and is being offered since Fall of 2023. Designed to meet the ABET as well as the CAE-CO requirements, the degree is created to meet the fast-growing demand for computer scientists with expertise in cybersecurity. The curriculum emphasizes hands-on coursework and experiential learning and equips students with state-of-the-art knowledge and skills to address security challenges related to data, software, hardware, connection, cyber systems, and cybersecurity threats impacting organizations and society.

WSU is also home for the Department of Defense-supported *VICEROY Northwest Institute for Cybersecurity Education and Research (CySER)*. The CySER Institute is established with a vision to train a diverse cybersecurity workforce for military or civilian careers in support of national defense. The institute has developed a successful training program that combines cybersecurity education with experiential learning realized via *mentored research, bi-weekly seminars, class projects, summer workshops, and internships*.

Since 2025, WSU also has an active CyberCorps Scholarship for Service (SFS) program that aims to educate and train the next generation of cybersecurity professionals for federal, state, local, and tribal government positions. The program is housed within the School of Electrical Engineering and Computer Science and leverages the recently launched Bachelor of Science in Cybersecurity degree program, the thriving VICEROY CySER Institute, and the extensive breadth of faculty expertise in cybersecurity, computer science, computer engineering, and electrical engineering. SFS scholars receive comprehensive hands-on training in security across the entire computing stack (hardware, systems, software, web) and they learn how emerging artificial intelligence and cryptography concepts can be used to build more resilient cyber systems. The program emphasizes training in six interrelated themes: (1) artificial intelligence and security, (2) cyber-physical systems security, (3) cryptography and post-quantum security, (4) software supply chain security, (5) hardware security, and (6) web security. In this presentation, we will share our experience from establishing the SFS program and training the first cohort of SFS scholars as well as our experience from running the VICEROY CySER program over the last five years.

Keywords: Artificial intelligence and security, cryptography and post-quantum security, cyber-physical systems security, hardware security, web security, software security, cyber education.

2026 CAE in Cybersecurity Symposium – Pittsburgh, Pennsylvania
Refereed Proceedings – CAE in Cyber Research (R) Track

A Publication of the CAE in Cybersecurity Community National Center

West Virginia University: Cyber research and defense ecosystem

[School Highlight Talk]

Matthew Valenti, West Virginia University, WV, Matthew.Valenti@mail.wvu.edu

Anurag Srivastava, West Virginia University, WV, Anurag.Srivastava@mail.wvu.edu

Extended Abstract

West Virginia University (WVU) is a nationally recognized hub for cybersecurity innovation, holding dual designations as a CAE-R and CAE-CD since 2006. As an R1 research institution, WVU supports an interdisciplinary cybersecurity ecosystem spanning engineering, computing, business, law, health sciences, and public policy. This collaborative environment addresses challenges in cyber defense, secure systems, critical infrastructure resilience, and trusted identity, while developing a highly trained cybersecurity workforce. WVU's CAE-R program is driven by a diverse faculty with expertise in areas including software and systems security, networks, hardware security, biometrics, AI security, robotics security, and cyber-physical resilience, supported by growing undergraduate and graduate programs such as the B.S. in Cybersecurity and NSF-funded workforce initiatives.

WVU has demonstrated strong research productivity and impact, with more than 100 cybersecurity publications in the past three years and substantial external funding from agencies such as NSF, DOE, DoD, DARPA, and NASA. Notable efforts include cyber-power resilience assessment capabilities, hardware-in-the-loop testbeds for securing AI-enabled systems, and secure bio-cryptosystems for trusted identity. This work is enabled by a robust research infrastructure that includes smart grid and cyber-physical security labs, a network operations center with AI-driven analytics, cyber ranges, and specialized facilities for biometrics and next-generation wireless security.

The university maintains long-standing partnerships with federal agencies, industry, and regional organizations, including a 20-year collaboration with NASA, participation in DoD and U.S. Cyber Command programs, and engagement with industry leaders. WVU also contributes to operational cybersecurity readiness through initiatives such as Exercise Locked Shields and the Cyber-Resilience Resource Center, which supports small businesses across the state. Through these integrated efforts in research, education, and outreach, WVU advances cybersecurity innovation while translating research into practice and workforce development.

Keywords: Critical Infrastructure Resilience, Trusted Identity, AI & Hardware Security.