


# Cybersecurity Curricular Guidelines

Matt Bishop<sup>1</sup>, Diana Burley<sup>2</sup>, Scott Buck<sup>3</sup>, Joseph J. Ekstrom<sup>4</sup>,  
Lynn Futcher<sup>5</sup>, David Gibson<sup>6</sup>, Elizabeth K. Hawthorne<sup>7</sup>, Siddharth Kaza<sup>8</sup>,  
Yair Levy<sup>9</sup>, Herbert Mattord<sup>10</sup>, and Allen Parrish<sup>11</sup>

<sup>1</sup> University of California at Davis, Davis, CA, USA  
mabishop@ucdavis.edu

<sup>2</sup> The George Washington University, Washington, DC, USA  
dburley@gwu.edu

<sup>3</sup> Intel Corp., Hillsboro, OR, USA  
scott.buck@intel.com

<sup>4</sup> Brigham Young University, Provo, UT, USA  
jekstrom@byu.edu

<sup>5</sup> Nelson Mandela Metropolitan University, Port Elizabeth, South Africa  
Lynn.Futcher@nmmu.ac.za

<sup>6</sup> United States Air Force Academy, Colorado Springs, CO, USA  
david.gibson@usafa.edu

<sup>7</sup> Union County College, Cranford, NJ, USA  
hawthorne@ucc.edu

<sup>8</sup> Towson University, Towson, MD, USA  
skaza@towson.edu

<sup>9</sup> Nova Southeastern University, Fort Lauderdale, FL, USA  
levyy@nova.edu

<sup>10</sup> Kennesaw State University, Kennesaw, GA, USA  
hmattord@kennesaw.edu

<sup>11</sup> United States Naval Academy, Annapolis, MD, USA  
aparrish@usna.edu

**Abstract.** The goal of the Joint Task Force on Cybersecurity Education is to develop comprehensive undergraduate curricular guidance in cybersecurity that will support future program development and associated educational efforts. This effort is a collaboration among the ACM, the IEEE Computer Society, the AIS Special Interest Group on Security and Privacy (SIGSEC), the IFIP WG 11.8, and the Cyber Education Project. In January 2017, the Joint Task Force released a draft of those guidelines. This paper describes the framework underlying the guidelines, examines one set of topics, and then places this work in the context of an exemplary curriculum on cybersecurity education.

**Keywords:** Cybersecurity education · Curricular guidance · CSEC2017

# 1 Introduction

Recent accelerated growth in the number and variety of computing security education academic, training, and certification programs has led to an increased interest in what a cybersecurity professional<sup>1</sup> should know and what skills they should have. Agreement has been uncommon; disagreement is the norm. There is no agreed-upon body of knowledge that such a professional should know, no agreement on a specific set of practices that a cybersecurity professional should have experience with, and no agreement on the competency levels of the desired skills.

This lack of commonality arises from the nature of cybersecurity. It encompasses many disciplines, and is used in many roles. As a result, the knowledge and skills that a cybersecurity professional will use varies depending upon the job description. The knowledge and skills of one who makes policy differs considerably from one who architects defenses for a given installation. Certainly, both need to know something of what the other does; but the depth of knowledge and set of skills will vary considerably. Underlying all knowledge and skills of cybersecurity professionals, though, is a core body of knowledge that all should know and experiences they should have, regardless of role.

With this in mind, the ACM, IEEE-CS, AIS SIGSEC, and IFIP WG 11.8 assembled a Joint Task Force on Cybersecurity Education.<sup>2</sup> The goal of this working group is to develop cybersecurity curricular guidelines for undergraduate programs that emphasize different areas of specialization. For example, the guidelines will provide a basis for a training institute's certification program, a university's degree program in cybersecurity, and for a business school's MBA program. The certification program would emphasize practice, the university's program both theory and critical thinking, and the business school would emphasize policymaking to support a company's business mission.

This approach avoids the conflicts that arise when one tries to prescribe a common body of knowledge that defines a "cybersecurity professional". More important is the considerable flexibility of this approach. One program might emphasize the role people play in cybersecurity, and so incorporate literature, sociology, psychology, and language classes into its program. Another might simply focus on how to design and implement a security-based network architecture, for example that protects medical records housed at a hospital. Both can draw from the guidelines. Both will emphasize aspects of the guidelines that provide their professionals (or prospective professionals) with the knowledge and skills they need. Both can decide to discount specific areas, but will be aware that they are doing so—and knowing what one does not know is a hallmark of a well-educated, knowledgeable, skilled person.

The goal of this paper is to present the framework and basis for the guidelines in their current form, and examples of the guidelines themselves and how they

---

<sup>1</sup> We use this term to mean anyone working in an occupation requiring her to protect data, cyberinfrastructure, or computing resources.

<sup>2</sup> <http://www.csec2017.org>.

might be used. The draft of these guidelines is called the CSEC2017, and is available for comment [2]. The final version is still under development.

## 2 Background

“Cybersecurity” is a widely used term that speaks to the security of systems and data but has many different definitions. Examples include “[p]revention of damage to, protection of, and restoration of computers, electronic communications systems, electronic communications services, wire communication, and electronic communication, including information contained therein, to ensure its availability, integrity, authentication, confidentiality, and nonrepudiation” [9]; “defensive methods used to detect and thwart would-be intruders” [13]; and “the organization and collection of resources, processes, and structures used to protect cyberspace and cyberspace-enabled systems from occurrences that misalign *de jure* from *de facto* property rights” [7]. Agresti [3] notes that four forces (rebranding, organizational imperative, cyberspace domain, and national defense priority) shape the definitions. Thus, any guidance aimed at cybersecurity must begin by defining that term.

To make clear the scope of the guidance, the Joint Task Force defined “cybersecurity” as:

A computing-based discipline involving technology, people, information, and processes to enable assured operations. It involves the creation, operation, analysis, and testing of secure computer systems. It is an interdisciplinary course of study, including aspects of law, policy, human factors, ethics, and risk management in the context of adversaries.<sup>3</sup>

Cybersecurity first arose as a technical issue when multiprocessing systems were developed; then, the question was how to keep processes from interfering with one another. Time-sharing raised issues of users interfering with one another, and theory and controls were developed to understand the problem and prevent or hinder compromise [6, 12, 15, 17]. As a discipline, cybersecurity first appeared in a 1970 report from a group chaired by Willis Ware [16]. The U.S. Air Force subsequently chartered a Computer Security Technology Planning Study [4], led by James P. Anderson; that 1972 study defined many basic concepts such as the Trojan horse. The theory of cybersecurity was developed [5, 8, 10, 11, 14]. As networking grew, the need for cybersecurity increased, and the Internet and the World Wide Web stretched the reach of attackers so much that even non-technical users were put at risk. Attackers became more ingenious, and defenses improved, and the attackers then improved. This cyber arms race continues to this day.

As the discipline of cybersecurity has grown, so have educational programs, professional training programs, and certification programs. Each emphasizes

---

<sup>3</sup> The CSEC2017 draft has the last words, “in the context of adversaries” at the end of the first sentence of the definition [2, p. 10]. That is a misprint.

aspects of cybersecurity in their own way. Cybersecurity jobs span the gamut from homes to small offices to governments to international corporations. The wide variety of jobs, and hence professionals, in this field has made defining an educational body of knowledge difficult, because each type of job requires different knowledge and skills. Thus, cybersecurity education programs should be based on core cybersecurity knowledge and skills. They should have a computing-based foundation, teach concepts that are applicable to a broad range of cybersecurity expertise, and emphasize ethical responsibilities and obligations. Finally, many programs will tailor their curriculum so their graduates can go into specialities that are in demand at the time.

The goal of the CSEC2017 guidelines is to provide a basis for developing such programs. It begins with a model that unifies these concepts and views them through different specialities, as well as the application of the knowledge, skills, and concepts.

### 3 Model

The model consists of four parts:

1. Knowledge areas, the basic organizing structure and core ideas;
2. Cross-cutting concepts, which span the knowledge areas;
3. Disciplinary lens, which provides views of the model based upon specific disciplines; and
4. Application areas, which help define the level of coverage for each knowledge areas.

In this section, we discuss each of these parts in detail.

#### 3.1 Knowledge Areas

The six *knowledge areas* define the subject matter of cybersecurity. They form a body of knowledge for practitioners, researchers, teachers, and others. They are composed of *knowledge units*, which describe the sets of topics and what students should know about each. Each unit also describes learning outcomes.

The knowledge areas meet three criteria:

1. The area is important for multiple disciplines;
2. The area provides a tool for understanding or exploring cybersecurity ideas; and
3. The material in the area can be learned in varying levels of detail and understanding over time.

The areas, and a brief description of what each encompasses, follows.

- **Data Security.** This knowledge area focuses on the protection of stationary and moving data. It requires an understanding of both algorithms and analysis, and deals with both the theory and application of these. Example units in this area are confidentiality, integrity, and cryptography.

- **Software Security.** This focuses on the design, development, implementation, deployment, maintenance, and operation of software that meets security requirements, both explicit (security) and implicit (robustness). It includes all types of assurance in software, reverse engineering, and analyzing and handling malware, as well as what is often termed “secure software”.
- **System Security.** This area speaks to the composition of components that make up a “system” such as a computer or the infrastructure supporting networks. Aspects of this area dealing with software focus on its integration and use as a component of a system rather than the security of the software (although that may affect how the software is handled). The supply chain, digital forensics, devices (hardware), authentication, access control, and cyberphysical systems fall into this area.
- **Human Security.** Protecting people’s data in the context of organizations (i.e., as employees) or personal life, and their privacy, is a critical task of cybersecurity that this area covers. It also includes security-related behaviors such as how people react to social engineering attacks, social engineering itself, and identity management.
- **Organizational Security.** The organizational security knowledge area deals with security in the context of organizations. The type, size, and function of the organizations are not constrained. An important element of this area is risk—what it is and how it can be mitigated in the context of the organization. Examples of other knowledge units are disaster recovery, business continuity, compliance, and security evaluations.
- **Societal Security.** The ubiquitousness of computers, networks, and devices that computers control makes cybersecurity a necessity in society. This knowledge area deals with those facets of cybersecurity that impact society as a whole. Example knowledge units are ethics, cyber law and crime, professional codes of conduct, intellectual property, and cultural constraints and controls on cybersecurity processes, procedures, and technologies.

There is some overlap among these areas, in the sense that a knowledge unit often can be put into more than one. A good example is the design of a library that reads and processes packets from a network. This can go in the software security knowledge area, because it is a question of assurance: how does the library ensure that the packets are protected and handled as required? It can also go into the system security area, as the library interface is used in the composition of components (that is, network to system). The above organization minimizes this overlap, and when it occurs suggests that the same knowledge unit would have two different views, each view based upon the knowledge area.

### 3.2 Cross-Cutting Concepts

These concepts provide a framework for making connections among the knowledge areas. They unify underlying ideas, and so help students understand the material in the areas regardless of the discipline the student encounters them in. The model defines the following cross-cutting concepts:

- **Confidentiality (C)** is a property defined by rules that control the spread of information. Such a set of rules may define who can and cannot access data or resources, for example.
- **Integrity (I)** is a property that describes the accuracy and trustworthiness of information. A key component of this property is assurance, which defines the evidence provided to convince the audience that the objects meet some desired level of accuracy or trustworthiness.
- **Availability (A)** is a property defined by rules describing when and in what manner data or resources can be accessed. Note that mere accessibility of an entity is not enough to make that entity available. If, for example, a network connection to a server does not meet the required quality of service (the property defining availability in this context), the entity may be accessible but not available.
- **Risk (R)** describes the exposure of the entity to threats. It is a product of the probability of the threat being realized and the damage incurred should the threat be realized.
- **Adversarial Thinking (AT)** is a manner of thinking in which one determines how threats can be realized. It requires understanding what threats will compromise the entity under consideration, and how to realize them.

As an example of why these are cross-cutting, consider “confidentiality.” It is a key component of data, system, and organizational security. In the guise of privacy, it is a component of human security. Through the combination of all these, it speaks to societal security. So the precise meaning of the term depends on the context in which it is used, as does its application. But the underlying concept of confidentiality, that of restricting access to something based on a set of requirements and a (possibly unstated) policy, cuts across these knowledge areas.

Each of these areas is broken down into units, and the units into topics. For example, the System Security knowledge area is broken down into units that include availability and secure system design, which in turn have the topics system availability, measures of availability, and attacks on availability; and security design principles, security architectures, trusted computing base, and security modes of operations, respectively.

### 3.3 Disciplinary Lens

While the knowledge areas are common to all of cybersecurity, the depth and approach that students and practitioners are expected to know varies depending on how they will use the knowledge. For example, a programmer needs to know something about the policies that control the data her program will use, as data in medical records must be handled very differently than data in a store’s inventory. Similarly, a policymaker needs to know that software cannot distinguish between the “good guys and gals” and the “bad guys and gals”, so backdoors in software and systems to aid law enforcement can be used by others, illicitly. The


programmer need not know how policy is made; the policymaker need not understand how backdoors can be implanted or exploited. Both need to understand enough of the other's world to see the consequences of their actions.

For curricular purposes, we use the disciplines as defined by the ACM [1]:

- **Computer Science (CS)** is the discipline of developing software, developing ways to use computers to solve problems, and indeed developing new ways to use computers. It includes theory and applications, and is a broad subject.
- **Computer Engineering (CE)** focuses on the design and implementation of computing devices. This requires understanding how the devices will be used, what software they must run, and in what environment they will be used.
- **Information Systems (IS)** deals with use of information processing technology in enterprises such as businesses, with an emphasis on the use of information that can be obtained from the use of the systems. It examines how to integrate that technology into the processes by which the enterprise attempts to meet its goals.
- **Information Technology (IT)**, like information systems, focuses on information technology but with an emphasis on “technology.” This encompasses what type of technology is appropriate for the needs of the organization, how the technology is to be deployed and maintained, and how to support both the technology and the users of that technology.
- **Software Engineering (SE)** is the discipline of defining, developing, implementing, testing, and maintaining software systems.
- **Mixed Disciplinary (MD)** programs contain elements chosen from the disciplines above.

### 3.4 Summary

A program can view the knowledge areas and cross-cutting concepts through the appropriate disciplinary lens to determine which concepts it should consider



**Fig. 1.** The relationships of the elements making up the thought model.

core, and which need to be touched only lightly upon. Figure 1 summarizes how these elements are put together.

## 4 Professional Practice

The CSEC2017 model for cybersecurity curricular guidelines are linked to professional practice through seven application areas. Workforce frameworks can then codify bodies of knowledge for their target audience by going from the application areas back to the model, and extracting both the core knowledge and cross-cutting concepts they deem appropriate, and view them through the disciplinary lenses appropriate for their audience.

The application areas are simply organizing frameworks to allow the definition of competency levels needed for each area. The content of the areas overlaps, as is expected; each area provides guidance for the depth of coverage needed for each idea.

The application areas are organized along the lines of the system and software life cycle as well as supporting areas. They are:

- **Public policy** is affected by several groups. Legislators and regulators make laws affecting the development, deployment, and use of computing. In the U.S., judges pass upon the constitutionality (legality) of the laws and of their application to specific cases and circumstances. Corporate managers (CEOs and members of the Board of Directors or similar entity) will also interact with public policy, either advising those who set it or ensuring that their computing technology and procedures comply with those policies. They must understand how these laws, regulations, and requirements will affect the use of the systems, how people interact with the systems, and most especially the risks that those rules reduce or increase. This suggests they know basics of the design of systems, so they understand what computing technology can—and, more critically, cannot—do, and be able to work out the budgetary and human costs of the rules.
- **Procurement** requires understanding how the systems will fit into, and advance, the work of the organization. This may involve changing aspects of the organization’s procedures to enable the systems to be useful. Risk management and business continuity issues come into play in this application area, as does an understanding of how people in the organization, and other stakeholders, will interact with the systems being procured. So this application area requires a knowledge of those areas, in addition to understanding concepts of assurance, infrastructure, and organizational, human, and social dynamics.
- **Management**, which refers to people, systems, and data in an organization, is guided by policies, both internal and public. Understanding compliance, business continuity, and recovery from attacks and other types of disasters is a part of management. Managers decide who has access to data and resources, what type of access, and how they may use that access; thus, they must understand identity and authorization management. As they will oversee, and be



responsible for, the effect of changes made to the system, understanding how both assurance and testing speak to the goals of the organization and of the mission of that particular system. They must also have a basic understanding of incident handling and recovery in order to deal with attacks.

- **IT security operations** are to keep the system secure, “secure” being defined by a set of requirements. This requires that operations personnel know how to translate those requirements into configurations, procedures, and their implementation. For example, the security infrastructure must ensure that identity management systems are installed, initialized, configured, and used properly. Validating that the requirements are properly implemented requires testing the infrastructure, systems, and procedures, and analyzing the results of those tests. In addition, the operators must be able to maintain the systems under both normal conditions and under abnormal conditions, such as during an attack.
- **Software development** begins with requirements for the software to meet. These requirements come from laws, regulations, policies, business plans, and societal and organizational constraints. So developers interpret these requirements in their design and implementation. The development must ensure the software is robust (“secure programming”), which means they must know how to determine which exceptions to handle and understand how to do so. The environment, users, and installers all must be taken into account when the software and interfaces are designed. Testing enables the developers to gather assurance evidence to verify the software system meets its requirements, and convince other stakeholders this is so.
- **Research** requires all researchers to know the basics of access control and availability, confidentiality, integrity, risk, and adversarial thinking. Cryptography is commonly used to supply the last two areas, so its basics are important to know. Beyond that, the specific area in which the research is being conducted defines what the researcher needs to know. For example, a researcher in cryptography should understand how it is used in practice in order to understand how the application affects the parameters of the cryptosystems; it is probably unnecessary to understand the proof of the HRU theorem and the associated results. But someone studying formal models of access controls would need to know the proof of the HRU theorem, and not the details of cryptography.
- **Enterprise architecture** is in some sense a capstone of the application areas, as it draws on all the other areas. Policy drives the architecture; the design of the architecture drives procurement, management, and operations. The architecture also affects the goals of the software, because the architecture includes the systems and infrastructure needed to keep the enterprise running smoothly. Enterprise architects must understand the policy, procurement, management and operations application areas, as well as elements from the area of software development.

## 5 Conclusion

CSEC2017 presents the basis for curricula. An institution desiring to have a specialization, major, or other course of study in cybersecurity may use this to design and implement their program. No single program will be able to cover all knowledge areas and cross-cutting concepts in full depth; instead, they should cover these broadly, and select specific aspects of those areas to explore in more depth. The selection will be based on the goals of the program, and the needs of the students attending the program and of the workforce that they expect, or will be expected to, join.

The inclusion of human, organizational, and societal knowledge areas emphasizes that cybersecurity is not a strictly technical discipline. The humanities and social sciences play a key role in cybersecurity. The connection with social sciences is clear, as those deal with society and organizations. The connection with humanities is equally important, because art, literature, languages, and other such subjects teach about the human condition and about people—and ultimately the goal of cybersecurity is to enable people to protect people, through the medium of guarding data and resources. Hence, including these two humanistically-oriented subjects in a cybersecurity program increases the likelihood that cybersecurity will protect the right people, data, or resources, in the right way, and at the right times.

Cybersecurity as a discipline is still maturing. Community input is important, and the Joint Task Force encourages comments, suggestions, and improvements. The 2016 International Security Education Workshop, held in June 2016 in Philadelphia, PA, USA provided one such avenue; a subsequent global stakeholder survey in late 2016 provided more input. The CSEC2017 web site, <http://www.csec2017.org>, provides a mechanism to submit comments. In this way, the curriculum guidance will meet the needs of the cybersecurity teaching, research, and practice workforces.

**Acknowledgements.** We gratefully acknowledge the valuable contributions of participants in our 15 community engagement efforts.

This material is based upon work supported by the National Science Foundation under Grant No. DGE-1623104, the National Security Agency’s CNAP Curriculum Development effort (RFI-2017-00022), the Education Board of the ACM, and Intel Corporation. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation, the National Security Agency, the ACM Education Board, or Intel Corporation.

## References

1. Computing curricula 2005: the overview report. Technical report, ACM, New York (2005). [http://www.acm.org/education/education/curric\\_vols/CC2005-March06Final.pdf](http://www.acm.org/education/education/curric_vols/CC2005-March06Final.pdf)

2. Cybersecurity curricula 2017: curriculum guidelines for undergraduate degree programs in cybersecurity. Technical report Draft version 0.5, ACM Joint Task Force on Cybersecurity Education (2017). <http://www.csec2017.org/csec2017-v-0-5>
3. Agresti, W.W.: The four forces shaping cybersecurity. *IEEE Comput.* **43**(2), 101–104 (2010)
4. Anderson, J.: Computer security technology planning study. Technical report ESD-TR-73-51, ESD/AFSC, Hanscom AFB, Bedford, MA, October 1972
5. Bell, D.E., LaPadula, L.J.: Secure computer system: Unified exposition and multics interpretation. Technical report MTR-2997 Rev. 1, The MITRE Corporation, Bedford, MA, USA, March 1976
6. Conway, R.W., Maxwell, W.L., Morgan, H.L.: On the implementation of security measures in information systems. *Commun. ACM* **15**(4), 211–220 (1972)
7. Craigen, D., Diakun-Thibault, N., Purse, R.: Defining cybersecurity. *Technol. Innov. Manage. Rev.* **4**(10), 13–21 (2014). <https://timreview.ca/article/835>
8. Denning, D.: A lattice model of secure information flow. *Commun. ACM* **19**(5), 236–243 (1976)
9. Dukes, C.W.: Committee on national security systems (CNSS) glossary. Technical report CNSSI No. 4009, Committee on National Security Systems, National Security Agency, Ft. George G. Meade, MD, USA. <https://cryptosmith.files.wordpress.com/2015/08/glossary-2015-cnss.pdf>
10. Graham, G.S., Denning, P.J.: Protection: principles and practice. In: *AFIPS Conference Proceedings: 1971 Fall Joint Computer Conference*, vol. 39, pp. 417–429. ACM, New York, November 1971
11. Harrison, M., Ruzzo, W., Ullman, J.: Protection in operating systems. *Commun. ACM* **19**(8), 461–471 (1976)
12. Hoffman, L.J.: The formulary model for flexible privacy and access controls. In: *AFIPS Conference Proceedings: 1972 Spring Joint Computer Conference*, vol. 40, pp. 587–601. ACM, New York, May 1972
13. Kemmerer, R.A.: Cybersecurity. In: *Proceedings of the 25th International Conference on Software Engineering*, pp. 1–11 (2003)
14. Lampson, B.W.: Protection. *ACM SIGOPS Operating Syst. Rev.* **8**(1), 18–24 (1974)
15. Saltzer, J.: Protection and the control of information sharing in multics. *Commun. ACM* **17**(7), 388–402 (1974)
16. Ware, W.: Security controls for computer systems: report of Defense Science Board Task Force on computer security. Technical report R609–1, Rand Corporation, Santa Monica, CA, February 1970
17. Weissman, C.: Security controls in the ADEPT-50 time-sharing system. In: *Proceedings of the 1969 Fall Joint Computer Conference*, pp. 119–133. ACM, New York, November 1969