



Scope for the Cyber Security Body of Knowledge

Awais Rashid, George Danezis, Howard Chivers, Emil Lupu and Andrew Martin

Version 2.0

November 10, 2017

<http://www.cybok.org>

Acknowledgements

We thank Makayla Lewis (for conducting the scoping research through community workshops, interviews and the online survey), Claudia Peersman (for implementing and applying relevant natural language analysis techniques to analyse the large amounts of textual materials used in the scoping phase) and Yvonne Rigby for managing the project excellently, especially in a condensed timescale for the scoping research.

We also thank everyone who contributed insights for the scoping research through the various engagement activities, which were invaluable in identifying the Scope. Special thanks are due to our Professional Advisory Board and International Academic Advisors for their insights and advice during the scoping work. We are also grateful to Mark Gondree and Ashley Podhradsky for including the CyBOK panel in the programme for the Advances in Security Education Workshop at USENIX Security Symposium 2017.

We are thankful for the comments received from members of the National Cyber Security Centre (NCSC) as well as Fred Piper and Steve Furnell during our progress review meetings with NCSC.

Last, but by no means least, we gratefully acknowledge the support from the National Cyber Security Programme for funding the research, and the interest shown by Cabinet Office and DCMS.

Executive Summary

Context

The National Cyber Security Strategy, published in 2016, set out a well-resourced plan to make Britain confident, capable and resilient in a fast-moving digital world. The approach acknowledged the importance of developing cyber skills within and across professions, engaging industry and allies. This need to address the cyber security skills gap is echoed by governments, researchers, educators and practitioners internationally. The efficient delivery of effective education and training programmes calls for the development of an accessible and internationally respected body of knowledge to which experts can contribute, maintaining coherence and currency.

The CyBOK Project

The Cyber Security Body of Knowledge (CyBOK) project aims to codify the foundational and generally recognised knowledge on cyber security. CyBOK is meant to be a *guide* to the body of knowledge—the knowledge that it codifies already exists in literature such as textbooks, academic research articles, technical reports, white papers and standards. The project's focus is, therefore, on mapping *established knowledge* and not fully replicating everything that has ever been written on the subject. Educational programmes ranging from secondary and undergraduate education through to post-graduate and continuing professional development programmes can then be developed on the basis of CyBOK.

Establishing the CyBOK Scope

Since the 1st of February 2017, the project team undertook a range of community consultations, both within the UK and internationally, through a series of different activities designed to gain as much input as possible and from as wide an audience as possible. In addition, analysis of a number of relevant texts, such as tables of contents of textbooks, calls for papers for conferences and symposia, standards, existing certification programmes, etc. was undertaken to complement the insights gained from the community consultations. The insights from these activities were synthesised to develop a Scope for CyBOK and 19 top-level Knowledge Areas (KAs) identified. The initial CyBOK Scope and KAs identified were made publicly available for community comments. While none of the 19 KAs needed to be removed or new ones added, the topics to be covered under each KA were refined based on input received. The 19 KAs are summarised at the end of this executive summary. The KAs are grouped into five broad categories. The grouping of KAs is provisional and may be subject to change. It should be read as a guide to structure, and not a prescription of content.

Next Steps in the Development of CyBOK

Internationally recognised experts will be invited to author detailed descriptions of each KA, which will be reviewed by a small panel of peer-reviewers before being made available for public consultation. As each KA description is finalised, it will be made available on the CyBOK web site. We aim to complete all KA descriptions by the end of July 2019. Alongside, learning pathways through CyBOK and exemplar curricula at different education levels will be developed.

Cyber security is a rapidly changing and evolving field. As such the CyBOK will never be 'finished' per se. Future iterations will need to be undertaken to ensure that the coverage remains up-to-date and the KAs reflect both current state of knowledge in cyber security and emerging needs. The inclusion of KAs such as Hardware Security and Cyber-Physical Systems Security in the current Scope reflects such emerging needs. Any future maintenance of CyBOK will need to ensure that, whilst not ignoring the needs of contemporary and legacy systems, the CyBOK scope also reflects key challenges arising from the increasing integration of technology – and hence cyber security – into the very fabric of our society.

The 19 Knowledge Areas (KAs) within CyBOK

Human, Organisational and Regulatory Aspects	
<i>Risk Management and Governance</i>	Security management systems and organisational security controls, including standards, best practices and approaches to risk assessment and mitigation.
<i>Law and Regulation</i>	International and national statutory and regulatory requirements, compliance obligations and security ethics, including data protection and developing doctrines on cyber warfare.
<i>Human Factors</i>	Usable security, social and behavioural factors impacting security, security culture and awareness as well as the impact of security controls on user behaviours.
<i>Privacy and Online Rights</i>	Techniques for protecting personal information, including communications, applications and inferences from databases and data processing. It also includes other systems supporting on-line rights touching upon censorship and circumvention, covertness, electronic elections and privacy in payment and identity systems.
Attacks and Defences	
<i>Malware and Attack Technologies</i>	Technical details of exploits and distributed malicious systems, together with associated discovery and analysis approaches.
<i>Adversarial Behaviours</i>	The motivations, behaviours and methods used by attackers, including malware supply chains, attack vectors and money transfers.
<i>Security Operations and Incident Management</i>	The configuration, operation and maintenance of secure systems including the detection of and response to security incidents and the collection and use of threat intelligence.
<i>Forensics</i>	The collection, analysis and reporting of digital evidence in support of incident or criminal events.
Systems Security	
<i>Cryptography</i>	Core primitives of cryptography as presently practised and emerging algorithms, techniques for analysis of these and the protocols which use them.
<i>Operating Systems and Virtualisation Security</i>	Operating systems protection mechanisms, implementing secure abstraction of hardware and sharing of resources, including isolation in multi-user systems, secure virtualisation and security in database systems.
<i>Distributed Systems Security</i>	Security mechanisms relating to larger scale coordinated distributed systems, including aspects of secure consensus, time, event systems, peer-to-peer systems, clouds, multi-tenant data centers, and distributed ledgers.
<i>Authentication, Authorisation and Accountability</i>	All aspects of identity management and authentication technologies, and architectures and tools to support authorisation and accountability in both isolated and distributed systems.
Software and Platform Security	
<i>Software Security</i>	Known categories of programming errors resulting in security bugs, and techniques for avoiding these errors - both through coding practice and improved language design, and tools, techniques and methods for detection of such errors in existing systems.
<i>Web and Mobile Security</i>	Issues related to web applications and services distributed across devices and frameworks, including the diverse programming paradigms and protection models.
<i>Secure Software Design and Development</i>	The application of security software engineering techniques in the whole systems development lifecycle resulting in software that is secure by default.
Infrastructure Security	
<i>Network Security</i>	Security aspects of networking and telecommunication protocols, including the security of routing, network security elements and specific cryptographic protocols used for network security.
<i>Hardware Security</i>	Security in the design, implementation, and deployment of general-purpose and specialist hardware, including trusted computing technologies and sources of randomness.
<i>Cyber-Physical Systems Security</i>	Security challenges in cyber-physical systems, such as IoT and industrial control systems, attacker models, safe-secure designs, security of large-scale infrastructures.
<i>Physical Layer Security</i>	Security concerns and limitations of the physical layer including aspects of radio frequency encodings and transmission techniques, unintended radiation, and interference.

Contents

1 Introduction	7
1.1 The CyBOK project	7
1.2 The Scope Document	7
1.3 Community Consultation on the CyBOK Scope	7
1.4 Providing Further Comments and Feedback	8
1.5 Next Steps in the Development of CyBOK	8
2 Scoping Research	8
2.1 Online survey	9
2.2 Analysis of relevant texts	9
2.3 Interviews with key experts	10
2.4 Community workshops	10
2.5 Call for position statements	11
3 Knowledge Areas	11
4 Human, Organisational and Regulatory Aspects	12
4.1 Risk Management & Governance	12
4.2 Human Factors	13
4.3 Privacy & Online Rights	13
4.4 Law and Regulation	14
5 Attacks and Defences	15
5.1 Malware and Attack Technologies	15
5.2 Adversarial Behaviours	16
5.3 Security Operations & Incident Management	17
5.4 Forensics	17
6 Systems Security	18
6.1 Cryptography	18
6.2 Operating Systems & Virtualisation Security	19
6.3 Distributed Systems Security	20
6.4 Authentication, Authorisation & Accountability (AAA)	21
7 Software and Platform Security	21
7.1 Software Security	21
7.2 Web & Mobile security	22
7.3 Secure Software Lifecycle	23

8 Infrastructure Security	24
8.1 Network Security	24
8.2 Hardware Security	24
8.3 Cyber-physical Systems Security	25
8.4 Physical Layer & Telecommunications Security	26

1 Introduction

1.1 The CyBOK project

Cyber security is recognised as an important element in curricula at all educational levels. However, the foundational knowledge upon which the field of cyber security is being developed is fragmented and, as a result, it can be difficult for both students and educators to map coherent paths of progression through the subject. The overall aim of the Cyber Security Body of Knowledge (CyBOK) project is to codify the foundation and generally recognised knowledge on cyber security following a broad community engagement within the UK and internationally. CyBOK will be a *guide* to the body of knowledge—the knowledge that it codifies already exists in literature such as textbooks, academic research articles, technical reports, white papers and standards. The project's focus is, therefore, on mapping *established knowledge* and not fully replicating everything that has ever been written on the subject.

1.2 The Scope Document

Since the 1st of February 2017, the CyBOK project team has undertaken an extensive exercise involving a mapping and analysis of relevant texts as well as a range of community consultations via workshops, an online survey, interviews and position papers. These activities are summarised in Section 2 and have provided an in-depth understanding of the community's collective view of the top-level Knowledge Areas (KAs) that should be in the Scope of CyBOK.

Following these consultations and various inputs, we have distilled 19 top-level KAs that will be in the Scope of the CyBOK. This document provides an overview of these top-level KAs and the sub-topics that should be covered under each KA. The KAs are discussed in detail in Section 3.

1.3 Community Consultation on the CyBOK Scope

CyBOK is aimed as a resource *for the community and by the community*. Throughout the development of CyBOK, inputs and reviews from the worldwide community of researchers, practitioners, educators and other stakeholders will be continually sought.

Version 1.0 of the Scope document was released on 19 September 2017 to seek community input and feedback on the Scope that had been distilled from the previous consultations. We invited members of the cyber security community and stakeholders from academia, industry, practice and professional organisations to comment on the CyBOK Scope. Comments were sought through consultation workshops, an online form as well as direct input through email or annotated copies of the Scope document.

Feedback was sought via the following key questions:

Question 1: *What are the strengths of the current CyBOK Scope?* Examples responses include but are not limited to: coverage in terms of breadth of scope and/or coverage of particular KAs. In a nutshell, we are interested in understanding what you appreciate about the proposed CyBOK Scope.

Question 2: *What are the improvements that can be made to the current CyBOK Scope?* Similar to Question 1, example responses include but are not limited to: comments on coverage in terms of breadth of scope and/or coverage of particular KAs. However, comments mustn't simply state that something is imperfect. We would like to receive proposals on how to improve upon any issues highlighted. Proposals may include suggested rewordings, additional topic areas to be covered and so on. In a nutshell, we are interested in understanding how the Scope may be improved through specific updates.

Question 3: *Are there other comments not covered by Questions 1 and 2?* We welcome other comments on the Scope document as colleagues may see fit. Similar to questions 1 and 2, we

request that these are constructive and make concrete proposals for improvement.

The CyBOK team reviewed and synthesised all comments and feedback received and updated the Scope as suitable. This document constitutes Version 2.0 of the Scope following these revisions.

1.4 Providing Further Comments and Feedback

We will continue to welcome comments on the CyBOK Scope to further refine the document. These can be sent by email to contact@cybok.org either as a free form text or an annotated PDF document. Comments can also be provided online at: <http://www.cybok.org/>.

1.5 Next Steps in the Development of CyBOK

The development of detailed KA descriptions will begin from the 1st of November 2017. We will also be undertaking research on more detailed knowledge dependencies across KAs within CyBOK and knowledge beyond what is captured in CyBOK. We will also be developing exemplar learning pathways through CyBOK.

The details of the methodology used to develop the CyBOK and various processes are publicly available:

<https://www.cybok.org/media/downloads/methodology-v1.1.pdf>

Development of each KA will be overseen by an editor, who will normally be a member of the Project Management Board¹. KAs will be assigned to an author (or authors) and will also have an expert review panel. This panel will include five international experts who will provide detailed scrutiny of the KA description and offer comments and feedback to the authors. The guidelines for authors invited to write the detailed KA descriptions are publicly available:

https://www.cybok.org/media/downloads/Brief_for_Authors_v_2.1.pdf

Each author will be invited to write a description of the knowledge areas, according to guidance given in the Author Guidelines using the Scoping Document as a guide and each iteration will be reviewed by the expert review panel. Once the content has been agreed and approved, it will pass to a public review phase. Comments from this public review phase will be synthesised and addressed by the KA authors. As each KA becomes stable, it will be published on the CyBOK web site. When all KAs are finalised, the collection will be published as a complete CyBOK on the web site.

2 Scoping Research

Since the 1st of February 2017, community consultations have been undertaken through a series of different activities, designed to gain as much input as possible and from as wide an audience as possible. The activities also provided a means to bring the project to the attention of the wider community. In addition, analysis of a range of relevant texts, such as tables of contents of textbooks, calls for papers for conferences and symposia, standards, existing certification programmes, etc. was undertaken to complement the insights gained from the community consultations. The various activities are summarised below:

¹Details of the project management structure and various boards are available on: <http://www.cybok.org/>

Data collection activity	Level of engagement
Online survey	44 responses received
Analysis of relevant texts	44 separate texts analysed
Interviews with key experts	10 interviews undertaken
Community workshops	11 workshops completed 106 attendees in total
Call for position statements	13 statements submitted
Panel at Advances in Security Education Workshop at USENIX Security Symposium 2017	Paper-based exercise with 28 attendees

There was a largely even distribution of representation from academia and practitioner organisations during most of the consultations:

Data Collection Activity	Academic (%)	Practitioner (%)
Online survey	51	49
Interviews with key experts	50	50
Community workshops	55	45
Position statements	62	38

2.1 Online survey

An online survey was published on the project website. It was launched at the UK Academic Centres of Excellence in Cyber Security Research conference on 28 June 2017 and remained open until 31 July 2017. The survey offered an opportunity for the community to provide views on the Scope of CyBOK and the Knowledge Areas to be covered by means of a series of open- and closed-ended questions. In addition to demographic data, the survey sought participants' views on topics such as: the cyber security knowledge areas that had been most important background knowledge in their career; key knowledge areas that ought to be covered in the CyBOK and those that should be out of scope; and topics that would be of most importance over the next 5 years. These and other questions were used to elicit participants' views on the CyBOK Scope.

2.2 Analysis of relevant texts

A diverse cross-section of different types of knowledge and texts was analysed. We used a variety of text mining techniques, such as Natural Language Processing (NLP) and automatic text clustering to cluster relevant topics and identify relationships between topics. Techniques utilised included semantic word cloud visualisations, Word Vectors, Ward clustering, K-means clustering and Latent Dirichlet Allocation (LDA). The documents analysed included:

- Categorisations, such as the ACM CCS taxonomy;
- Certifications, such as CISSP and the IISP Skills Framework;
- Calls for Papers such as IEEE Symposium on Security and Privacy, USENIX Symposium on Usable Privacy and Security;
- Existing curricula, such as ACM Computer science curriculum, work of the ACM Joint Task Force on Cyber Security Education;
- Standards, such as BS ISO-IEC 27032 2021, NIST IR 7292;
- Table of contents of various textbooks.

We note that the above are examples of some of the texts analysed and do not represent an exhaustive list.

2.3 Interviews with key experts

Interviews with key experts were conducted nationally and internationally to elicit views on the Scope of CyBOK and the knowledge areas to be covered. The interviews were undertaken during July, August and into early September 2017. These semi-structured interviews were used to gather verbal data from 10 key experts using an interview guide. Lines of enquiry were pursued within the interviews to follow up on interesting or unexpected avenues that emerged.

The experts were recruited based on their international standing in research or practice. The interviewees were not merely technical experts in cyber security. Cyber security also includes topics such as governance, regulation, risk and law, and experts with knowledge of these topics were also interviewed.

2.4 Community workshops

A set of participatory workshops was designed that brought together over 100 attendees from industry and academia to discuss the knowledge areas that should be included in CyBOK in a collaborative and creative environment. Some workshops were dedicated to consultation with academia and others to consultation with practitioners. A subset also included representatives from both academia and practitioner communities. Invitations to the workshops were distributed via various established networks and communication channels. They were also publicised on the CyBOK project website and registrations were managed through the EventBrite system.

The workshops schedule ran from 05 June 2017 to 21 July 2017 at various locations around the UK including Lancaster, London, Edinburgh, Nottingham, Birmingham and Cardiff to ensure maximum coverage. Initially only 5 workshops were planned but, due to demand from the community, the number of workshops was increased to 11, all of which were well attended with an average attendance of 10.

The workshops were based on a *shopping trolley* metaphor whereby participants were encouraged to think about what they believe are the key KAs to be included by putting each KA into one of the four *supermarket* areas:

- *In the trolley* – knowledge areas to be included;
- *On the shopper's heart* – knowledge areas that are of interest to participants but not necessarily to be included;
- *On the shelf* – knowledge areas to be discussed further;
- *In the bin* – knowledge areas deemed out of Scope.

This sorting exercise was followed by a *15 items or less* task during which participants were asked to sort the 'in the trolley' KAs into groups of top-level and sub-level KAs. This workshop design allowed for small group discussion on where knowledge areas should be best placed and why. It also led to sub-topics within knowledge areas to be identified.

In addition to these workshops, consultations were also held at the Higher Education Academy Conference in Liverpool, UK in April 2017 and the Cyber Security Professionals Conference in York, UK in May 2017. A panel discussion was also organised at the Advances in Security Education Workshop at the USENIX Security Symposium in Vancouver, BC, Canada in August 2017 and participants' views on importance of particular topics sought via a paper-based exercise.

2.5 Call for position statements

A call for short position statements was launched to invite contributions from the international community towards the identification of knowledge areas. The call opened on 11 May 2017 and closed on 30 June 2017. The position statements covered both technical and social aspects of cyber security and included representation from both academics and practitioners.

3 Knowledge Areas

The data generated from each of the activities in Section 2 was initially analysed separately to identify key trends and topics. The various analyses were then collated and synthesised to identify commonalities. The commonalities were, in turn, used to help define the Scope and the KAs that comprise the Scope. In total, 19 top-level KAs were distilled, grouped into five broad categories, as shown in Figure 3. We note that other possible categorisations of these KAs may be equally valid. Nor are the categories necessarily orthogonal.

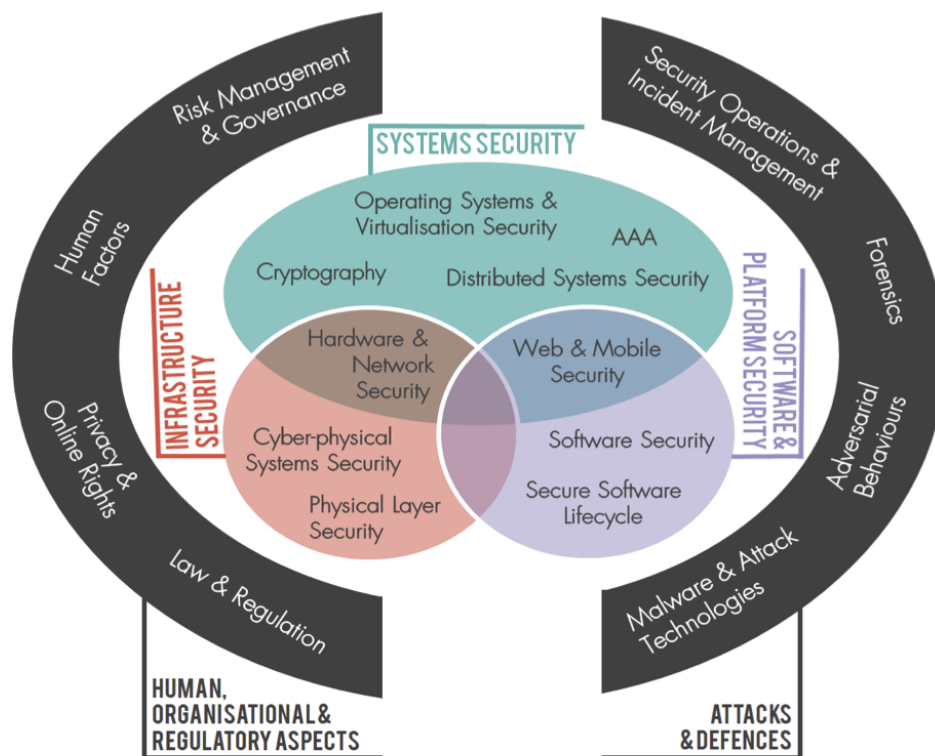


Figure 1: The 19 Knowledge Areas (KAs) in the CyBOK Scope

We next describe these KAs in more detail. Each KA description includes a brief overview followed by a short summary describing the extent of the KA. This summary is indicative and may be subject to change in the detailed drafting process. However, any substantial divergence by the authors will need approval from the Project Management Board, which may in turn consult the International Academic Advisors and the Professional Advisory Board. For each KA, we also note its key dependencies on other KAs in CyBOK as well as on knowledge external to the CyBOK Scope.

4 Human, Organisational and Regulatory Aspects

4.1 Risk Management & Governance

Security management systems and organisational security controls, including standards, best practices and approaches to risk assessment and mitigation.

This KA covers current best practices in security and risk management systems. Specifically, it focuses on how cyber security risk is socially constructed when humans, technologies and organisational processes intersect. It will thus address security issues at an organisational level and how to identify and implement relevant security policies, standards, procedures and guidelines to mitigate potential cyber security risks.

The KA will discuss and contrast both bottom-up (i.e., system-centric) and top-down (risk management) approaches to risk assessment. It will also review the challenges of unravelling cyber security in complex organisational settings where risk perceptions of individuals or groups can vary significantly. The impact of individual or group perceptions on risk decision-making and how this, in turn, impacts the overall risk posture of an organisation will be discussed. Challenges of aligning cyber security with organisational aims and high-level strategic objectives will also be covered along with issues such as gap between technical measures to mitigate risks and the information required for board-level risk decision-making.

The KA will also cover techniques for threat assessment, asset audit, vulnerability analysis and impact assessment as well as those for risk mitigation, such as business continuity management, disaster planning and recovery. The merits and demerits of quantitative and qualitative (as well as pseudo-quantitative techniques) will be contrasted. In this context, specific techniques and frameworks, for instance, Hermann's security metrics, ISO 27005, NIST SP 800-30, CRAMM, FIRM, OCTAVE, etc., may be discussed as illustrative examples. Issues of objective and traceable risk measurements will be captured along with the challenges of deriving such measurements.

Security management systems and standards will also be discussed. This will include fundamental information security principles such as Confidentiality, Integrity and Availability and techniques for formulation and implementation of organisational security policies and controls to implement such principles in practice. Relevant security management standards such as PCI, NIST Cybersecurity Framework, FIPS and ISO 27000 series will also be discussed.

Challenges of security governance and decision-making including those arising from business, economic and organisational factors will be discussed. Techniques and strategies for board-level oversight of cyber security and risk management at an organisational level will be described. Topics of security and risk culture within organisations will also be covered. Specific challenges of risk management and governance in particular contexts such as cyber-physical systems may also be discussed insofar as they differ from typical enterprise environments. However, any detailed discussion of methodologies ought to be deferred to the relevant KA.

Depends on KAs: Human Factors (individual and group perceptions and behaviours impacting risk); Law and Regulation (impact of regulatory landscape on organisational policies and security/risk management practices); Adversarial Behaviours (understanding of threat actors); Malware and Attack Technologies (risks arising from particular types of attacks).

Depends on External Knowledge: General work on risk management and organisational/business studies.

4.2 Human Factors

Usable security, social and behavioural factors impacting security, security culture and awareness as well as impact of security controls on user behaviours.

This KA covers security issues, challenges and opportunities arising from the intersection of humans and technologies. It focuses on the fact that cyber security behaviour is shaped by individual and group processes and, equally, technology is made vulnerable and is exploited by the individual. Taking such an embedded view of cyber security makes it possible to encapsulate the behavioural and technological aspects of existence and security in the digital world and provokes new thinking about established distinctions between such concepts as online/offline, attacker/insider, risk/protection, etc.

The KA will discuss the need to consider cyber security in its wider socio-technical context, highlighting how security is often a secondary activity, that is, users' focus tends to be on the primary task. Design of security systems needs to consider this secondary nature of security as an activity. The KA will cover the substantial body of literature on usable security, highlighting how design of security systems and controls hampers usage of a system and how this in turn shapes users' behaviours with regards to security. Well-known issues with usability of encryption tools, passwords and disconnect between security policies and users' work practices will be discussed. Usability issues arising from design of security mechanisms in emergent technologies such as cloud and mobile platforms will also be discussed. Topics such as users' mental models of security and how these impact their security decisions and behaviours will be reviewed.

The discussion will also cover software developers as users of libraries and application programming interfaces (APIs) and the security issues arising from lack of usability of such libraries and APIs.

The KA will also capture extant knowledge on psychological and behavioural studies on victims of cybercrime and cyber attack. It will also discuss techniques and approaches to cultivate a cyber security culture and improve cyber security awareness and education amongst users and across organisations. This will include the growing body of literature on shifting the focus from *humans as the weakest link* to *humans as a security asset and resource* and the impact of such a shift on the design of security systems, policies and processes.

Issues of user and data exploitation arising from the participatory data economy will also be covered and a discussion of ethical considerations around user data collection, monitoring and sharing will be included.

Depends on KAs: Adversarial Behaviours (for impact of adversarial behaviours on users/victims); Risk Management & Governance (for organisational and governance aspects impacting users); Law and Regulation (for the impact arising from the legal and regulatory landscape on individuals).

Depends on External Knowledge: Social and behavioural sciences methodologies for studying human and organisational aspects.

4.3 Privacy & Online Rights

Techniques for protecting personal information, including communications, applications and inferences from databases and data processing. It also includes other systems supporting online rights touching upon censorship and circumvention, covertness, electronic elections and privacy in payment and identity systems.

This KA describes knowledge about the engineering of systems that are mindful of user privacy, and technologies that protect user and data privacy, through hiding it (confidentiality), through allowing users to control how their information is processed (control), and through providing visibility and transparency into how personal information is used (transparency). It complements the discussion of purely regulatory aspects of privacy protection, covered in the Law and Regulation KA, and discusses

how those rules can be backed up by strong technological underpinnings. Other technologies relating to basic human rights such as election systems, and censorship circumvention are also covered.

We expect that classic and advanced privacy enhancing technologies would be covered: those include techniques for network security, including anonymous communications and their traffic analysis; secure channels tuned to privacy protection, such as OTR; and systems implementing cryptographic protections, such as PIR, presence, address book privacy. The KA also covers a discussion of modern censorship technologies, and options for covert communication to protect privacy and circumvent censorship, including steganography and steganalysis. Real-world architectures for mass surveillance can be described.

In terms of data privacy, the KA covers discussions of privacy preserving data processing. The field of inference control in statistical databases, modern definitions of private statistics, such as differential privacy (DP) and DP mechanisms are discussed – as well as deployed and proposed systems implementing such mechanisms. Advanced topics such as privacy friendly statistics collection, and privacy-friendly machine learning are also covered. Electronic election systems are also reviewed.

Mechanisms, that implement control and transparency are also covered: for example the use of selective disclosure techniques as part of identity systems, etc; privacy features of authentication protocols; user-centric identity; and their relevance to electronic cash systems, or other privacy-preserving tokens. Pure cryptographic primitives are covered by the Cryptography KA, but privacy-specific systems and applications can be covered here.

Abuse prevention techniques, to prevent anonymity or privacy being used as a vector for attack, can be mapped. Regulatory aspects of privacy are not covered here, but other aspects such as the economics of privacy, discrimination, profiling, and technologies to support data protection (e.g., computer readable privacy policies and audit), etc can be covered in this KA. Attack vectors against privacy, including profiling, side-channels in web applications or fingerprinting of devices can be covered.

Depends on KAs: Cryptography (for privacy-oriented crypto); Law and Regulation (for censorship and elections, and data protection); Authentication, Authorisation & Accountability (AAA) (for privacy in authentication and identity); Human Factors (for usability issues pertaining to privacy mechanisms); Secure Software Lifecycle (for engineering privacy into the design of systems).

Depends on External Knowledge: Sociological and behavioural methodologies and studies.

4.4 Law and Regulation

International and national statutory and regulatory requirements, compliance obligations including data protection, and developing doctrines on cyber warfare.

This KA describes statutory and regulatory requirements which are relevant to security policies or influence how computer systems are designed, operated and used. The content provides an international perspective, contrasting US and European approaches and noting significant differences in other territories.

The KA will address regulations that require, specify or influence cyber security. The protection of personal data and data privacy regulations, crimes committed against or using computers, and monitoring and interception of information are important concerns in most jurisdictions. In addition to unauthorised access etc., criminal use of computers extends to prohibited content such as child abuse or terrorist material, and monitoring and interception includes permissible security monitoring and the resulting requirements for user consent, as well as requirements to support law enforcement by providing access to data. The KA will cover such topics as well as additional legal constraints and obligations associated with intellectual property, and with conducting business via electronic means. Where such topics are conditioned by other legislation, such as Human Rights, the impact of this context on cyber security compliance will also be described.

Transborder issues will also be discussed including data protection and privacy issues, multilateral agreements such as mutual co-operation in criminal investigation and the attribution of crimes or acts of cyber warfare.

Emerging legislation may be enacted in one territory but avoided or debated in others, and such trends will be separately described. Examples include security standards for critical infrastructure, the sharing of security incident information, responsibility for content by search or social media providers, and in some territories censorship and elections.

Depends on KAs: Malware and Attack Technologies (for jurisdictional issues arising from, e.g., botnets); Adversarial Behaviours (transborder crime); Privacy & Online Rights (legal frameworks for protecting users' rights and data); Risk Management & Governance (impact of law on governance and organisational approaches to risk management).

Depends on External Knowledge: Background understanding of legal systems and legislative processes.

5 Attacks and Defences

5.1 Malware and Attack Technologies

Technical details of exploits and distributed malicious systems, together with associated discovery and analysis approaches.

This KA will describe the technical aspects of computer attacks and malicious software and hardware. The higher-level aspects relating to economics, strategy and cyber crime are covered by the associated KA on Adversarial Behaviours. In terms of malware, it will cover the standard taxonomy of malicious software, its operation and function, such as viruses, trojans, botnets, etc.

It will discuss the architecture of historical and modern malware, such as the separation between injection, command & control, and payload, as well as techniques for long term persistence such as root kits. Techniques for weaponising exploits to be used as part of malware will be covered along with modern and historical techniques to establish command and control in a survivable manner and tactics to avoid detection, including polymorphism and techniques for engineering malware including kits.

This KA will also discuss software and hardware specifically geared towards detecting and eliminating malware and attacks, beyond the engineering aspects of systems: such as anti-virus, system integrity protections, and host-based intrusion detection research. Techniques and tactics for analysing malware, including malware deploying countermeasures will also be reviewed including reverse-engineering, emulation, unpacking and packing and laboratory procedures for the safe handling of malware samples and their containment. Further, modern security operations aiming to disrupt and take over networks of malicious software will be discussed – including botnet infiltration and take down techniques, possibly illustrated in their technical details by existing case studies in the scientific literature.

Specific and exotic forms of malware will also be discussed, including hardware implants, complex malware relating to cyber physical systems (such as Stuxnet), and malware affecting office applications, browsers, mobile platforms or IoT where there is a necessary knowledge base. Besides unambiguously malicious systems, this KA will also covers the grey zone of adware and unexpected software installers and browser toolbars. Specifics about the monetisation, payload, ecology and cybercrime aspect of malicious software are covered in the Adversarial Behaviours KA.

Depends on KAs: Adversarial Behaviours (for attacker behaviours and economics of cyber crime); Several: Network Security, Operating Systems & Virtualisation Security, Distributed Systems Secu-

rity, Web & Mobile security and Cyber-physical Systems Security (for relevant security architectures targeted and exploited by malware).

Depends on External Knowledge: Operating systems; Programming language models and data structures; graph theory.

5.2 Adversarial Behaviours

The motivations, behaviours and methods used by attackers, including malware supply chains, attack vectors and money transfers.

This KA describes the characteristics of adversarial threat agents, their use of the network to anonymously support supply chains and attacks, and the strategies and vectors used in attacks.

Adversarial agents include both internal and external agents whose actions range from inadvertent or unstructured events to sophisticated hacking by criminal organisations or nation states. The KA provides a taxonomy of agent types together with a description of their typical capability, intent and likely targets.

Most adversaries develop and practice deliberate deception; ways in which deceptions are developed and maintained by different adversaries are described and used to set the context for the specific technologies, threat actions, and models described in this KA.

Such agents use anonymising network technology, such as onion routing and proxies, to support the supply of malware and related technology and to sell products such as user credentials, credit cards and data resulting from attacks. Transactions and the receipt of money from victims are facilitated by anonymous or semi-anonymous financial instruments, such as digital currencies, moneygrams, and virtual credit cards. The KA will cover extant literature that has studied such financial instruments and transactions and wider socio-economic considerations that influence adversarial behaviour. Economics of attacks, and relationship between providers and users of different attack technologies will be discussed along with new cyber crime models, e.g., crime-as-a-service (CAAS).

The KA will also cover models used to understand and analyse adversarial behaviours including attack trees, kill chains, and other widely used attacker models from literature, together with examples of how attacks are mapped to such models. Other types of attack such as denial of service, internal fraud or the exploitation of open-source intelligence, which may have different and distinctive templates, will also be described.

Attack vectors are an important part of adversary tradecraft and include phishing, drive-by websites and logic bombs as well as external electronic attacks. The command and control of a realised attack may also distinguish adversaries, both the control framework employed and how the attacker avoids attribution. Activities such as data and information gathering that underpin such attack vectors will also be covered, for example, the use of open-source intelligence (OSINT) and data from online social media.

Depends on KAs: Human Factors (impact of adversarial tactics on users/victims); Network Security (reconnaissance and exploitation of network); Security Operations & Incident Management (network monitoring and intrusion detection); Privacy & Online Rights (exploitation of OSINT); Law and Regulation (cross-border crime and legal frameworks).

Depends on External Knowledge: Crime science; criminology.

5.3 Security Operations & Incident Management

The configuration, operation and maintenance of secure systems including the detection of and response to security incidents and the collection and use of threat intelligence.

This KA focuses on technical aspects of operational security: system management, situational awareness, security monitoring and incident management. Human and organisational factors which are often integrated with operational security in practice are described in other KAs, namely Human Factors, Risk Management & Governance and Adversarial Behaviours. While physical security (i.e., of buildings, humans and infrastructures) is important, it is not in the scope of the KA description.

The KA will discuss system security functions such as managing the lifecycle of system components via configuration and change management, malware protection, and the provision of backup and recovery services. Security policies and their implementation and management, ranging from network segmentation to user access control, will also be covered.

The KA will also cover security situational awareness, that is, strategic input to audits and assessments by gathering information about potential threats and vulnerabilities. External sources such as threat intelligence services and security alerts and advisories will also be discussed along with how to supplement these by proactive organisational information gathering via honeypots and penetration tests.

Security monitoring to identify immediate threats, as well as providing a capability to analyse long term trends in incidents and user behaviour, will be discussed. Technical monitoring to consolidate and analyse inputs from system events and intrusion alerts will also be covered along with a discussion of how user behaviours, such as accesses or transactions, may be monitored to identify suspicious or fraudulent activity. Approaches to the integration of event, awareness, and security management information in Security Information and Event Management (SIEM) systems will also be described.

The KA will also cover Incident Management, including topics such as analysis and containment of the security incident, recovering operational capability and reporting of any lessons learned. Beyond methods of technical diagnosis, aspects such as team, management and communication will also be discussed.

Depends on KAs: Risk Management & Governance (Risk management approaches underpinning security operations and business continuity management); Human Factors (vetting, education and training, behaviour); Authentication, Authorisation & Accountability (AAA) (user access control and accountability aspects of monitoring and logging network data).

Depends on External Knowledge: Management science, including crisis response, team working, training, and decision making. Physical security.

5.4 Forensics

The collection, analysis and reporting of digital evidence in support of incident or criminal events.

This KA is concerned with the acquisition, analysis and reporting of digital evidence concerned with a security event or crime, to a standard which will allow the resulting evidence to be presented in court. The KA includes evidential requirements, gathering evidence, the forensic process and associated tools, core digital artefacts, and domain-specific issues.

Evidential requirements include the preservation of evidence, the traceability of results, and the integrity and repeatability of any analysis. Evidence gathering includes the established seizure and imaging process and also extends to situations where the baseline processes are not feasible, such as smartphones, volatile data, and selective recovery from network storage. Related process issues such as triage and search and discovery techniques will also be explored, as will approaches to

assurance of the tools used in the forensic process. Presentation and reporting aspects specific to forensic evidence will be described.

Core digital artifacts include disk and file systems, date and time records, operating system components, and evidence resulting from network activity. The KA coverage will extend beyond the base cases to include important modern variations, for example, data storage may include non-magnetic storage, document management systems, and distributed file systems.

Application domains of evidential importance or particular technical difficulty will be covered, for example, the detection of document or image manipulation and the processing of images for event detection or biometric recognition, location tracking, recovering evidence from embedded systems, and e-discovery. Anti-forensics will also be included, particularly the problem of encryption, and also the detection and analysis of other approaches to obfuscate forensic records.

Depends on KAs: Security Operations & Incident Management (forensics relating to incident management); Malware and Attack Technologies (Malware Analysis).

Depends on External Knowledge: Low-level functioning of computers and operating systems and of data representation in computers; Knowledge of criminal legal systems and procedures. General reporting and presentation skills.

6 Systems Security

6.1 Cryptography

Core primitives of cryptography as presently practised and emerging algorithms, techniques for analysis of these and the protocols which use them.

This KA describes the cryptographic theory, primitives, advanced constructions and proof techniques at the intersection of mathematics and computer science.

In terms of theory it covers definitions of security, such as unconditional information theoretic or complexity theoretic notions, against which the security of cryptography schemes is usually established. Specific important hard mathematical problems, forming the basis for public key or symmetric key cryptosystems can be discussed, with more advanced mathematical details referenced appropriately. Theoretical results such as those relating to one-way functions, their existence and necessity are touched upon. Primitives may be separated into symmetric key techniques, including the design of block ciphers, hash functions and stream ciphers; the reductions amongst them, as well as their modes of operation, with links to appropriate standards where applicable. Advanced modes of operation, such as constructions for authenticated encryption are discussed. Public key primitives include constructions for key exchange, encryption and digital signature schemes.

Modern definitions of security and proof techniques are presented: including provable security, game based definitions and models such as UC or reactive security. Families of protocols widely studied within cryptography are also covered. Those include zero-knowledge protocols, and succinct arguments of knowledge both from a theoretical standpoint as well as practical instantiations. Two-party or multi-party computation frameworks and their proof systems are also discussed, including those based on garbled circuits and secret sharing. Advanced primitives such as hash chains and Merkle Trees, homomorphic encryption, aggregatable signatures, attribute-based encryption and credentials, identity-based encryption, group signatures, ORAM, Private information Retrieval, and Oblivious Transfer, etc are also mapped. Their applications within specific systems are left to other KAs. Recent work on quantum algorithms and computing, post-quantum secure cryptosystems and primitives, as well as modern theoretical results related to obfuscation are discussed, with appropriate references to further material.

On the engineering end, this KA includes a discussion of secure implementation of cryptography primitives, implementation specific attacks (such as those using side-channels) and the state of the art in terms of the verification and validation of cryptographic software; a discussion of cryptographic acceleration and hardware for key management; as well as the necessity for key management and public key infrastructures (with details potentially left to the Network Security KA). All cryptographic background necessary for describing key deployed protocols including TLS, IPSec, SSH, Tor, should be described, while the security context in which these protocols operate will be provided within the Network Security and Privacy & Online Rights KAs.

Depends on KAs: Privacy & Online Rights and Network Security (for specific implementation and deployment of cryptographic protocols).

Depends on External Knowledge: Number theory; Information theory; Computer systems architecture (efficient implementation of bit operations, etc).

6.2 Operating Systems & Virtualisation Security

Operating systems protection mechanisms, implementing secure abstraction of hardware and sharing of resources, including isolation in multi-user systems, secure virtualisation and security in database systems.

This knowledge area introduces the principles and control mechanisms for ensuring security at the Operating System level. Starting with an introduction of basic principles: mediation, least privilege, Saltzer & Schroeder, this KA introduces system resources and constructs (CPU, memory, file systems I/O, processes, threads, etc.), and related protection mechanisms (privilege levels, memory and address protection, paging and segmentation, file system). It will present the basics of access control models and mechanisms, reference monitors, permissions, rings, capability based models and their use to protect objects inside the operating system. Particular emphasis will be put on separation/isolation and containment. Microkernels will be briefly introduced. Whilst the Authentication, Authorisation & Accountability (AAA) knowledge area introduces the basics of mandatory access control policies, this KA will present their concrete implementation in Trusted Operating Systems, TrustedBSD, SELinux, etc. Information flow control models (IFC and DIFC) will be presented as well as their implementation in operating system design and architecture. Starting from the integrity models described in the Authentication, Authorisation & Accountability (AAA) KA, and from the descriptions given in Hardware Security (hardware roots of trust), techniques for verifying and maintaining Operating System integrity will be introduced, such as secure boot, file system integrity e.g., (TripWire) etc. Standards and assurance in Operating Systems will be covered as well as the limitations of assurance.

Database security will extend the access control models presented in the particular context of Databases and also address specific issues such as reliability and integrity in database systems, consistency, queries and query aggregations, authenticated index structures, managing and querying encrypted data and introduce inference aspects covered in the Privacy & Online Rights KA to which it will refer. Issues of secure deletion from file systems and databases will also be covered.

Building upon the earlier introduction of microkernels this knowledge area will now present them in more detail and how they can be used to reduce the Trusted Computing Base. This aspect will also cover secure microkernel architectures and formally verified microkernels such as SEL4.

Resource virtualisation and emulation of BIOS, CPU, MMU memory and devices will then be presented leading to the notion of emulating entire virtual machines, type I (hypervisors) and type II virtual machine monitors (VMM) and admin VMs. Threat models for virtualised systems will be described including attacks on tenant OS, attack on hypervisor, cross-VM attacks as well as lower level attacks, e.g., on BIOS, SMM, DMA. Security architectures and protection models will be presented including VM-Introspection, use of Admin VMs, using Hypervisors for protection, protecting the hy-

pervisor itself (including through hardware assisted solutions), and using static and dynamic roots of trust (thereby linking into the Hardware Security KA).

Specifics of OS architectures and protection mechanisms for resource constrained devices embedded Operating Systems and RTOS as well mixed hardware software implementations and associated tradeoffs will also be presented.

Depends on KAs: Authentication, Authorisation & Accountability (AAA) (access control and integrity models); Hardware Security (roots of trust); Distributed Systems Security (utilisation of security architectures and hypervisors in cloud).

Depends on External Knowledge: Operating systems; Databases.

6.3 Distributed Systems Security

Security mechanisms relating to larger scale coordinated distributed systems, including aspects of secure consensus, time, event systems, peer-to-peer systems, clouds, multi-tenant data centers, and distributed ledgers.

This knowledge area comprises security issues linked to distribution, remote invocation, distributed services, and distributed architectures (middleware, peer-to-peer systems, distributed ledgers etc.). It also covers their use in security. Core principles and their security implications such as coordination, consensus, byzantine agreements and fault tolerance, voting, fairness etc. will be introduced. The protocols and other security aspects associated with distributed services will be covered including: distributed naming, distributed directories, e.g., LDAP, and service registries, e.g., UDDI, Hypercat, time services (including threats and implications of compromise), service discovery, membership management. Aspects associated with the security of remote invocation systems not covered elsewhere will also be described, for instance, RPCs, ORBs (and security services), Web-Service security and distributed file systems.

The security of other middleware architectures will be described and in particular the security aspects of messaging systems, e.g., MQ, rabbitMQ, event-based architectures and in particular publish-subscribe architectures (both topic-based and content-based), and stream processing systems.

Cloud security architectures and, in particular, aspects not directly covered under the Operating Systems & Virtualisation Security KA will be presented here. Types of cloud architectures SaaS, PaaS, IaaS and principles for cloud security architectures: secure isolation, data protection, secure deletion, monitoring and auditing will be presented. Security architectural patterns and guidelines including NIST Security, CSA and ENISA will be described. Compliance and reporting will also be discussed.

Security architectures of peer-to-peer (P2P) networks will be described including threats such as DDoS, Query Flooding and Poisoning. Vulnerabilities of distributed hash-tables and fairness aspects will also be covered. Trust Management and reputation in P2P systems will be described. Peer-to-peer network protocols, anonymous P2P communication and onion routing will also be covered.

This KA will also cover blockchains and distributed ledgers again both from the point of view of their operation and security as well as their use for security. Basic principles of operation of cryptocurrencies and distributed ledgers will be presented: proof of work, consensus layer, privacy and pseudonymity. Permissioned and permission-less ledgers will be discussed, as well as security aspects linked to denial of service, wallet management and side chains. Application of distributed ledgers, e.g., for non-repudiation, smart contracts, and crime related aspects will also be discussed.

Depends on KAs: Network Security (various issues pertaining to secure networking underpinning distributed systems); Operating Systems & Virtualisation Security (for secure isolation underpinning cloud and general operating system security issues).

Depends on External Knowledge: Distributed systems and programming.

6.4 Authentication, Authorisation & Accountability (AAA)

All aspects of identity management and authentication technologies, and architectures and tools to support authorisation and accountability in both isolated and distributed systems.

Starting with user account management and user enrolment this KA will include all techniques for end user identification and authentication including biometrics, multi-factor, continuous and behavioural authentication. The design of authentication protocols will then be described starting with symmetric key based protocols and their use in wired and wireless settings. The basics of protocol analysis, e.g., with BAN logic and protocol analysis tools (e.g., proverif, cryptoverif, f*, or easycrypt) also ought to be included. Protocols such as Kerberos will be described extending their scope to multiple domains. Similarly, authentication protocols based on asymmetric cryptography will be introduced together with their supporting infrastructure: certificates, PKI, pitfalls and governance of PKI infrastructure. The presentation will generalise to attribute management, attribute certificates, separately and when combined with identity. Technologies for federated identity management and authentication, e.g., RADIUS, SAML, Shibboleth, OAuth, etc. will be included in the description both as illustrative examples and in terms of their overall architecture. Trust management aspects and principles for trust negotiation will be introduced.

Authorisation aspects will build upon the knowledge of reference monitors and permissions and introduce principles for access control including mandatory, discretionary and policy based models. Mandatory models will cover both models for confidentiality, e.g., Bell-Lapadula and models for integrity, e.g., Biba, Clark & Wilson, Brewer & Nash. The RBAC family of models, its protection profiles, foundations (Named Protection Domains) and variations will be covered together with practical aspects of its implementations, e.g., role provisioning, mining and engineering. Policy based models, e.g., XACML, languages and logic for distributed authorisation (SecPAL, SDSI, TAOS, etc.) will be described leading the presentation to attribute-based distributed access control frameworks and the architecture of cross-domain authorisation models.

Accountability aspects will cover at least logging of user authentications, accesses and actions in single and distributed settings, use of applications and privileges. Log management will cover Log generation and storage, protection, analysis, retention and disposal. How to format, record, query, correlate and aggregate log information will be discussed along with integration with SIEM software.

Depends on KAs: Cryptography (entropy, one-way functions, protocol analysis); Human Factors (usability issues of AAA); Network Security (protocols); Operating Systems & Virtualisation Security (access control models implementation); Distributed Systems Security (distributed access control and cross-domain security issues); Security Operations & Incident Management (Log management and integration of SIEM software).

Depends on External Knowledge: Fundamentals of operating systems, distributed systems, networks and database systems.

7 Software and Platform Security

7.1 Software Security

Known categories of programming errors resulting in security bugs, and techniques for avoiding these errors – both through coding practice and improved programming language design, and tools, techniques and methods for detection of such errors in existing systems.

This KA will discuss the programming practices that lead to security bugs and the factors underpinning these, such as the programming models and type systems in use or the (unintended) misuse of particular language features or application programming interfaces (APIs). Vulnerabilities arising from

security bugs such as exposure of private information and man-in-the-middle attacks on supposedly secure communications will be reviewed along with information leakage via side channels and timing channels. Large-scale studies of software security bugs will be synthesised and key challenges for cultivating secure programming practices summarised.

This will be followed by a discussion of approaches, tools and techniques for supporting programmers during software development. Examples include type systems, verification techniques and automated tools for static and dynamic analysis. Programming language design issues will be discussed and design and implementation of safe languages covered. This may also include challenges arising from particular environments and coverage of secure programming models for mobile and cloud environments as well as programmable networks.

Specific techniques for improving software security will also be covered. Examples include secure runtime environments and protection mechanisms such as ASLR, Canaries, Control Flow Integrity and Automated Software Diversity techniques. Techniques for analysing and verifying that a software system upholds all desired security properties at runtime will also be discussed.

Depends on KAs: Malware and Attack Technologies (using bugs); Secure Software Lifecycle (for methodologies, such as code review); Security Operations & Incident Management (for security testing such as penetration tests); Human Factors (usability issues of APIs and programming models).

Depends on External Knowledge: Fundamentals of programming: abstraction, composition, programming models, type systems; Software verification.

7.2 Web & Mobile security

Issues related to web applications and services distributed across devices and frameworks, including the diverse programming paradigms and protection models.

This KA complements Operating Systems & Virtualisation Security, discussing the security systems around, and security issues of, mature mobile platforms such as smart phones as well as web technologies, including web browsers as platforms and web applications.

The commonality between those platforms includes a focus on isolation of applications based on provenance and origin; a security model based on permissions; the mobile code in the form of scripting; a security model around application stores and their procedures. Mobile platforms in particular are the subject of physical attacks to access privileged functions, and are composed of open and closed parts, including a physical layer processor that may come under attack. Their operating systems, while based on standard technologies, are geared towards isolation between apps rather than a multi-user setting, and sandboxing, virtual machines and type systems are deployed to preserve such isolation. Questions of trusted paths in and out are key, and also relate to issues of seeking user permission to perform operation to protect either security or privacy. Side channels may allow information to flow against policy, and platform weakness allow escaping sandboxes – with equivalent mitigation having been proposed. The KA will discuss these various topics.

This will be followed by a discussion of how web browsers are increasingly using similar techniques to maintain isolation, supported by a commodity operating system. Additional challenges associated with browsers will also be covered such as: their increased attack surface; the ubiquity of mobile code (such as javascript security models, and attacks); cross origin functionality, etc.

Besides the browser as a platform, key attacks and protection techniques have evolved relating specifically to web applications, and leveraging specifically the web. Injection attacks, cross site scripting, and other key attacks documented by OWASP and others will be discussed; including their appropriate mitigations. The architecture of web applications and micro-services will also be discussed, with a focus on vulnerabilities and their mitigation, including attacks that aim to escalate privileges on the server side. Web application firewalls, and other generic protection techniques will

be described. Authentication protocols specifically designed for web or mobile deployments will be reviewed, including technical issues around single sign on and identity systems – with higher level issues left for the Authentication, Authorisation & Accountability (AAA) KA. Key web application areas, such as social networking security and privacy, technical attacks on banking security, phishing, will be described when related to the technicalities of web or mobile security, with higher level concerns (such as the economics of cybercrime) left to other KAs.

Depends on KAs: Operating Systems & Virtualisation Security (secure operating systems, file systems and databases); Network Security (network architectures for web and mobile security); Distributed Systems Security (secure distributed systems and cloud platforms underpinning web and mobile platforms); Adversarial Behaviours (economics of cybercrime).

Depends on External Knowledge: Operating systems; Databases.

7.3 Secure Software Lifecycle

The application of security software engineering techniques in the whole systems development lifecycle resulting in software that is secure by default.

This KA describes how security is incorporated into the software engineering process, how the quality of the process and the resulting products may be assured, together with specialised technical approaches used in the development of secure software.

This KA will build on existing knowledge of practices in the standard software development lifecycle that are designed to progressively eliminate defects throughout the product development. It will focus on how a secure software engineering process essentially follows the same pattern while introducing additional security practices at each stage in the lifecycle. For example, architectural design requires consideration of security services such as authentication, software development needs code review and testing for known security defects, while deployment may require signed software. Secure system and software development processes such as NIST SP 800-160 and SDLC, together with security in Enterprise Architectures, will be discussed and used to motivate and illustrate security practices at each stage of the lifecycle. Topics such as patch development and management and decommissioning will also be discussed.

The KA will also cover assurance mechanisms, i.e., how the quality of the engineering process and of the resulting systems may be assured by using a certified assessment process, by benchmarking against good practice or maturity models, or by complying with best practice requirements which may be mandatory in some domains. Issues such as confidence in the assurance mechanism beyond compliance and its impact on decision-making will be discussed.

Specialised technical processes in security engineering will be covered where they represent established best practice. They may include formal approaches to requirements and system design, requirements abuse cases, attack modelling, automated code analysis for security defects, fuzzing, and hardware-software co-design. The KA will also cover incentive models whereby organisations promote third party testing of deployed products by offering bug bounties, bug auctions, competitions, or other inducements.

Cultural aspects of secure software development will also be discussed, such as cultures of development teams and how individuals or teams learn about security and keep their knowledge contextualised and up-to-date. Emerging scientific insights and best practices in this regard will be reviewed. The changing nature of software developers and hence, challenges of secure software engineering, whereby individuals develop and deploy mobile and web apps to potentially millions of users worldwide, will be discussed. Opportunities and challenges arising from socialisation of software development through forums such as StackOverflow will be covered.

Depends on KAs: Software Security (secure programming); Security Operations & Incident Management (penetration testing); Hardware Security (roots of trust for hardware-software co-design).

Depends on External Knowledge: General knowledge of software engineering lifecycles and techniques as covered in the Software Engineering Body of Knowledge (SWEBOK).

8 Infrastructure Security

8.1 Network Security

Security aspects of networking and telecommunication protocols, including the security of routing, network security elements and specific cryptographic protocols used for network security.

This knowledge area comprises all aspects relating to the security of wired and wireless networks above the physical layer. Different network technologies will be covered: IP networks (focus on V4 but also ought to cover main issues for IP V6), Wireless (802.11) and mobile networks (GSM, SS7, 3G, 4G, 5G). Threats, vulnerabilities and protection mechanisms at each network layer will be discussed as well as in virtualised networks. In particular, at layer 2 it is anticipated to cover MAC attacks in both wired and wireless networks, MAC layer misbehaviour, as well as VLAN hopping, etc. ARP and RARP aspects of security will be covered before moving onto the security of routing protocols that will be covered in detail in both wired (e.g., BGP) and wireless networks (e.g., MANET, sensor networks). Denial of service attacks and countermeasures, e.g., IP traceback, encapsulation, tunnelling, NAT and VPN and their security concerns will be covered as well as onion routing aspects. IPSec and IPv6 security will be described before moving onto the application layer protocols and services. The KA will also review security concerns in transport and application level protocols (TCP, DHCP, DNS) before describing secured versions e.g., SSH, TLS, DNSSEC.

The KA will also include descriptions of network defence tools including network perimeter protection, intrusion detection and prevention systems, anomaly detection and network monitoring, packet filters, stateful filtering, application gateways and firewall architectures. Security Operations & Incident Management are covered in a different KA.

Network resilience aspects including network management techniques for reconfiguration and recovery and adaptive architectures for network resilience, for instance, those based on a moving target defence approach, will be presented. Business continuity planning and recovery are covered in the Risk Management & Governance KA, whilst the lifecycle and recovery of system components is covered in the Security Operations & Incident Management KA. The use of Software Defined Network for security as well as the security of SDN itself will also be covered here. Finally, knowledge in this area should also cover aspects of integrity of network equipment itself and overall integrity of network data.

Depends on KAs: Authentication, Authorisation & Accountability (AAA) (various); Security Operations & Incident Management (network monitoring); Cryptography (encryption of data in motion); Physical Layer & Telecommunications Security (secure channels).

Depends on External Knowledge: General knowledge of networking and operating systems.

8.2 Hardware Security

Security in the design, implementation, and deployment of general-purpose and specialist hardware, including trusted computing technologies and sources of randomness.

This KA covers all aspects related to hardware-enabled and hardware-accelerated security in computer systems as well as the security of the hardware itself. Starting from identifying the role that hardware plays in the security of computer systems the content will cover physical tamper resistance and design principles for tamper resistant processors at different security levels, physical security aspects of magnetic media, storage and signal security. Moving from physical design to VLSI design,

integrated circuit design and then to board design, the KA will address, at each level, both threats and countermeasures. Hardware Trojans will be presented as well as techniques to detect and prevent them. Electromagnetic Analysis, Differential Power analysis, fault injection and differential fault analysis, timing attacks and other side channel attacks will be presented alongside prevention mechanisms such as random execution, fault tolerant architecture, unpredictable delay. Hardware implementation of cryptographic algorithms as well as hardware acceleration of cryptographic implementations will be included. Techniques for Intellectual Property protection in hardware designs, hardware reverse engineering and hardware obfuscation are also in scope.

Hardware sources of entropy and Physically Unclonable Functions will lead the way into a more general presentation of hardware roots of trust, secure co-processors, hardware security modules (HSMs). Secure processor architectures, trusted execution environments (TrustZone, SGXj, TEE Protection Profile), and trusted platform modules will be described in some detail leading into a presentation of secure boot and attestation techniques. Software attestation techniques will be presented covering both attestation with hardware roots of trust and software attestation as well as hybrid solutions of software attestation with secure hardware. NX Flags and hardware support for control flow integrity will also be covered. Trusted interfaces, DMA attacks and issues linked to the security of peripheral devices will be presented.

Depends on KAs: Cryptography (hardware implementation of cryptographic protocols), Authentication, Authorisation & Accountability (AAA) (for isolation models).

Depends on External Knowledge: Computer systems architecture; VLSI design.

8.3 Cyber-physical Systems Security

Security challenges in cyber-physical systems, such as IoT and industrial control systems, attacker models, safe-secure designs, security of large-scale infrastructures.

Cyber-physical systems (CPS) involve networked computational devices that utilise sensors and actuators to monitor and control a physical process. Example CPS include Internet of Things (IoT) based smart environments and critical infrastructures such as power grids, energy, water and manufacturing systems. This KA will discuss the unique security challenges arising from this integration of networked computational devices and control systems. Security issues such as those posed by shared CPS fabrics where potentially thousands of nodes may be deployed and used by a number of stakeholders to provide a multitude of services will also be discussed. Issues arising from the resource-constrained nature of devices in CPS environments will also be covered and relevant techniques that reduce the overhead, for instance, in terms of memory usage or computational power summarised.

Protocols specific to CPS, e.g., Modbus/TCP, EtherNet/IP, BACnet, Zigbee, WirelessHART, DNP3, etc. and their vulnerabilities will be covered.

The KA will also cover attacker models for CPS, the types of attacks that can be realised and defence mechanisms such as vulnerability scanners and intrusion detection systems that account for the specific properties (e.g., safety) and constraints (e.g., real-time operations) of a CPS. Different types of intrusion detection mechanisms, e.g., those that monitor communication channels, hardware or application properties over time, will be discussed.

The KA will also contrast the extant literature on known possible attacks and mitigation strategies for key CPS domains such as automotives, industrial control & SCADA (supervisory control and data acquisition) systems and IoT. Issues arising from the interaction between security and safety will be discussed and relevant work on safe-secure designs for CPS reviewed.

Depends on KAs: Operating Systems & Virtualisation Security (use of virtualised environments in CPS settings and cloud-based processing of IoT data); Network Security (network architectures for CPS), Hardware Security (trusted computing).

Depends on External Knowledge: Control systems; Safety-critical systems; Dependability.

8.4 Physical Layer & Telecommunications Security

Security concerns and limitations of the physical layer including aspects of radio frequency encodings and transmission techniques, unintended radiation, and interference.

This KA deals with security issues and mechanisms that are intimately associated with the interface between the low-level physical aspects of electromagnetic transmission and equipment, and the higher level telecommunication protocols or security systems.

With respect to communications security, the KA will cover: security related aspects of modulation and encoding; jamming and jamming resistant technologies, such as hopping, burst, spread spectrum, or meteor scatter communications. Interference and interference-free technologies will also be discussed along with aspects of acoustic signals when applicable. Classic RF related electronic warfare techniques will be covered, including direction finding, localisation, transmitter identification, and jamming.

The KA will also cover unintended emanations that could be used to compromise security systems. It will discuss the RF boundaries of security systems, TEMPEST and other compromising emanations, such as visual and acoustic; and techniques and standards, both software and hardware, to minimise those and evaluate RF related security.

At a higher level, protocols and systems relying on physical medium characteristics will be summarised. These include distance bounding protocols, as well as secure global positioning systems (such as Galileo). Applications of RF related techniques to secure sonar or radar, and jamming of those are also within scope. Finally, the engineering of physical layer hardware to resist attacks, and also generally hardware and software to ensure emanation security – and the associated standards – will be discussed.

Signalling and physical aspects of deployed telephony systems, such as POTS and mobile (GSM, 3G, 4G, 5G) will be covered. Security models, attacks and standards associated with telephony and data transmission networks (incl. SS7), as well as modern data modulation techniques from modems to fiber optic installations will also be reviewed.

Depends on KAs: Network Security (interface between physical layer security and higher-level protocols).

Depends on External Knowledge: Signal processing; Theory (RF, identification).