

Introduction

There is an increasing interest in research and development in information security. It is an area where active research is being conducted contemporaneously with application of the technology. For example, Internet and related information highway research and development are occurring while new applications are being added. Consequently, many research results are contained in technical reports that are available primarily within certain professional communities. The research results and applications reports that have appeared in the open literature (professional journals and conference proceedings) are scattered in several different sources. In spite of this, there aren't many new interdisciplinary books in this area. Those that are coming out focus primarily on application issues or address a narrow range of topics, such as computer viruses.

In assembling this collection of interdisciplinary essays, we have had a twofold objective: first, to provide a comprehensive summary of the results of the research, development, and application experience in information security up to this point, and second, to point toward directions for future work. We have attempted to obtain a balance among various viewpoints — some essays are research-oriented, some are from the producer's viewpoint, some from the consumer's perspective. We believe that this diversity of viewpoints provides a richness generally not available in one book. The diversity is evident in conflicts and contradictions among the essays. The opinions expressed are solely those of the individual authors and do not represent any organization. We have not tried to reconcile the differences. They reflect the dynamic nature of the subject matter.

To the greatest extent possible, we have tried to collect a coherent set of essays. The early essays provide background for those that follow. However, there is not a linear progression. Nor is there 100 percent uniformity in terminology. Concerning terminology, in some of the essays we have adopted a convention to identify words that are used in the context of information security in ways that differ from their common English usage. These "reserved words" are set in a different typeface. We hope that the change in typeface will remind you of the special usage without being too distracting.

Next we give brief summaries of all essays included in this collection. In Essay 1, "What Is There to Worry About? An Introduction to the Computer Security Problem," Don Brinkley and Roger Schell provide an over-

view of the vulnerabilities and threats to information security in computer systems. The essay begins with a historical presentation, contrasting the computer security problem with communication security problems. Next, it describes four broad areas of computer-related threats: theft of computational resources, disruption of computational services, unauthorized information disclosure, and unauthorized information modification. Classes of information-related threats are described, and examples of each are provided. These classes are *inadvertent human error*, *user irresponsibility*, *direct probing*, *probing with an artifice*, *direct penetration*, and *subversion of security mechanism*. The roles of Trojan horses, viruses, worms, bombs, and other kinds of malicious software are described and examples provided.

Essay 2, “Concepts and Terminology for Computer Security,” also by Don Brinkley and Roger Schell, provides an introduction to many of the concepts and terms that are most important in gaining an understanding of information security. It focuses on techniques for achieving access control within computer systems and networks. The essay begins by defining what is meant by information security and describing why it is important to constrain the definition to protection that can be meaningfully provided with a significant degree of assurance within computer systems. The theory of information security — the reference monitor concept — is introduced next through an analogy with information security concepts from the world of people and sensitive documents. Next, the essay further develops the presentation of the theory by introducing concepts and terms related to the security policy. Distinctions between discretionary and nondiscretionary security policies are provided, and supporting policies are introduced. Techniques for building a secure system based on the principles of the theory are presented, along with methods for usefully verifying the security of a system. The security kernel is presented as a useful, high-assurance realization of the reference monitor concept, and the principles behind the process of designing and implementing one from scratch are discussed. Improvements to the security of an existing operating system that are feasible, as well as fundamental limitations on those improvements, are described next. Finally, the reference monitor concept is applied to networks, and cryptography and access control are shown to be useful partners.

Without appropriate management, it is impossible to maintain security in a system or network of more than minimum complexity. In Essay 3, “A Philosophy of Security Management,” David Bailey discusses security management of complex systems, including the scope of the security manager’s role and the conflicting pressures that must be balanced. Bailey ends by discussing a strategy for the security manager that has been used successfully on a large local network.

Essay 4, “Malicious Software,” by Marshall Abrams and Harold Podell, discusses the threats, vulnerabilities, risk, effects, and countermeasures concerned with viruses, worms, and other forms of malicious software. It defines several common types of malicious software but avoids narrow semantic distinctions. The possible attacks include unauthorized modification of data and software (including the operating system) and unauthorized utilization of resources, often resulting in denial of service. The essay considers viruses in stand-alone computers as well as network vulnerabilities. Lessons learned from the study of various attacks are presented as general points for the future, including detection, reduction, recommendations, and legal remedies.

There are multiple views of corporate (enterprise) computing, each with its own metaphors and terms of reference. The different views incorporate different levels of abstraction, in which details are suppressed to concentrate attention on the issues important to the particular observer. Essay 5, “Abstraction and Refinement of Layered Security Policy,” by Marshall Abrams and David Bailey, examines these different metaphors with respect to the enterprise security policy. The result is a layered policy in which each main layer relates to one of the system metaphors, and the policy described for a lower level of detail is an implementation of the policy at a higher level. The layered view of policy helps system designers, managers, and users understand the rationale for security policy at the lowest levels of abstraction, because the relationship of the low-level policy to the enterprise information policy is clear.

The Trusted Computer System Evaluation Criteria (TCSEC) provide the basis for evaluating the effectiveness of security controls built into computer systems. In Essay 6, “Evaluation Criteria for Trusted Systems,” Roger Schell and Donald Brinkley summarize the definition and requirements of the TCSEC used to classify systems into seven hierarchical classes of enhanced security protection. This essay also summarizes the history, technical foundations, and basic security requirements of the TCSEC. These criteria have for a number of years been used in specifying security requirements during acquisition of products and systems, guiding the design and development of trusted systems, and evaluating systems used to process sensitive information.

In Essay 7, “Information Security Policy,” Ingrid Olson and Marshall Abrams discuss information security policy for automated information systems (AISs), focusing on information control and dissemination. Information security policy addresses such issues as

- disclosure, integrity, and availability concerns;
- who may access what information in what manner;

- basis on which the access decision is made (for example, user characteristic such as nationality or group affinity, or some external condition such as time or status);
- maximized sharing versus least privilege;
- separation of duties;
- who controls and who owns the information; and
- authority issues.

This essay discusses some of the aspects that must be considered when developing an information security policy for a given organization.

In Essay 8, “Formal Methods and Models,” James Williams and Marshall Abrams discuss how the motivation for using formal methods in the context of trusted system development stems primarily from their ability to provide precision, consistency, and added assurance during the elaboration of security requirements across different development stages. The subject matter of formal models and specifications can be illustrated by looking at the various kinds of security attributes and requirements that have turned up in published security models. Examples discussed in this essay relate to nondisclosure policies, data integrity policies, and user-controlled policies. This essay concludes with a discussion of technical and methodological issues relating to the effective use of formal methods.

Essay 9, “Rule-Set Modeling of a Trusted Computer System,” by Leonard LaPadula, describes a new approach to formal modeling of a trusted computer system. A finite-state machine models the access operations of the trusted computer system, while a separate rule set expresses the system’s trust policies. A powerful feature of this approach is its ability to fit several widely differing trust policies easily within the same model. LaPadula shows how this approach to modeling relates to general ideas of access control, and relates this approach to the implementation of real systems by connecting the rule set of the model to the operations of a Unix System V system. The trust policies demonstrated in the rule set of the model include the mandatory access control policy of Unix System V/MLS, a version of the Clark-Wilson integrity policy, and two supporting policies that implement roles.

Essay 10, “Representative Organizations That Participate in Open Systems Security Standards Development,” by Harold Podell, presents an introduction to representative organizations that participate in open systems security standards development required for interoperability and security of business and government computer network communications. International agreements include a commitment of the standards organizations to support open systems security standards to achieve “brand-independent” network configurations and interfaces. Four interrelated issues provide a basis for interpretation of current trends in the development of selected open systems standards:

1. the importance of security standards as an economic issue to support international electronic commerce;
2. a conceptual view of open systems security standards relationships;
3. a brief overview of the committee structure of the International Organization for Standardization (ISO), the International Electrotechnical Commission (IEC), the International Telephone and Telegraph Consultative Committee (CCITT), and selected national and regional organizations; and
4. an overview of important security standards committees' activities.

Essay 11, "Penetration Testing," by Clark Weissman, introduces flawless penetration testing as a requirement for high-rated secure systems — those rated above B1 based on the Trusted Computer System Evaluation Criteria (TCSEC). Penetration testing is a form of stress testing, which exposes weaknesses — that is, flaws — in the trusted computing base (TCB). This essay describes the Flaw Hypothesis Methodology (FHM), the earliest comprehensive and widely used method for conducting penetration testing; reviews motivation for penetration testing and penetration test planning; defines a flaw as a demonstrated unspecified capability that can be exploited to violate security policy; and provides an overview of the FHM and its analogy to a heuristic-based strategy game. Ten most productive ways to generate hypothetical flaws are described as part of the method, as are ways to confirm them. A review of the results and representative generic flaws discovered over the past 20 years is presented. The essay concludes with speculations on future methods of penetration analysis using formal methods: mathematically specified design, theorems, and proofs of correctness of the design.

In Essay 12, "Evaluation Issues," Marshall Abrams and Harold Podell present an introduction to evaluation issues in the US and European Community (EC) to illustrate the two schools of thought. Following development of draft national and regional criteria, the US, Canada, and the EC are working on Common Criteria (CC). The authors compare the proposed evaluation approaches in the hope that, in the international process of developing the Common Criteria, there will be a convergence to assist the multinational producers of secure products and systems in evaluation by different national entities. Interoperability of information technology and IT security to support electronic commerce depends, in part, on an acceptance of evaluated products in different countries and regions. Therefore, an important aspect of evaluation is the development of international agreements for reciprocity of evaluations of secure products.

Essay 13, "Supporting Policies and Functions," by Marshall Abrams and Harold Podell, observes that the major policy objective — to protect information assets against specific harm — usually requires additional

policies and functions for support and implementation. This essay discusses supporting policies and functions drawn from the TCSEC, the supporting “Rainbow” series, and the ITSEC.

Essay 14, “Security Engineering,” by Marshall Abrams, Harold Podell, and Dan Gambel, is concerned with trusted system integration and/or development to meet multilevel security (MLS) and operational requirements. It addresses technical issues such as how to combine products securely, TCB alternatives, and typical security engineering phases. It also addresses the management concerns of certification and accreditation.

Essay 15, “Cryptography,” which was written by Marshall Abrams and Harold Podell, discusses cryptographic protection of information confidentiality and integrity as that information passes from one point in space-time to another. More recent uses of cryptography, such as authentication and nonrepudiation, are also discussed.

The essay begins with an introduction of these ideas, including some basic examples, then proceeds to the definition of a cryptographic system, making the distinction between conventional key or symmetric key schemes and public key or asymmetric key schemes. The authors present some classical examples beginning with Julius Caesar. Both substitution and permutation ciphers are included, as well as a word about their weaknesses. The Data Encryption Standard (DES) serves as an example of a product cipher whose strength derives simply from repeated applications of both permutations and substitutions.

The essay then turns to public key schemes or systems. A public key system can be used by anyone to encrypt a message for a given recipient, but only that recipient can decrypt it. Although there are many proposed in the open literature and three have been widely implemented, the essay focuses on the most popular system — RSA. RSA (Rivest, Shamir, and Adleman) is a widely used public key system whose strength lies in the difficulty of factoring certain large numbers.

A discussion of public key management is followed by an introduction to public key and conventional key management issues. The authors also discuss authentication and integrity issues that are associated with conventional key systems. In addition, link encryption and end-to-end encryption are described and contrasted. The essay’s final topic is the integration of computer and communications security.

Essay 16, “Local Area Networks,” by Marshall Abrams and Harold Podell, addresses local area network (LAN) communications security. LANs are introduced as providing

1. a private communications facility,
2. services over a relatively limited geographic area,
3. a high data rate for computer communications, and
4. common access to a wide range of devices and services.

LANs share many security problems and approaches for their solutions with point-to-point conventional communications systems. In addition, LANs have some unique problems of their own:

1. universal data availability,
2. passive and active wiretap threats,
3. end-to-end access control, and
4. security group control.

Countermeasures include physical protection and separation by physical, logical, and encryption methods. Trusted Network Interface Units, encryption, and key distribution are also discussed. An example is discussed to illustrate different aspects of LAN security. The example is a composite of several existing product features, selected to demonstrate the use of encryption for confidentiality, and trusted system technology for local area networks.

Essay 17, "Internet Privacy Enhanced Mail," by Stephen Kent, presents Privacy Enhanced Mail (PEM) as consisting of extensions to existing message processing software plus a key management infrastructure. These combine to provide users with a facility in which message confidentiality, authenticity, and integrity can be effected. PEM is compatible with RFC-822 (Request for Comments¹) message processing conventions and is transparent to SMTP (Simple Mail Transfer Protocol) mail relays. PEM uses symmetric cryptography — for example, the Data Encryption Standard (DES) — to provide (optional) encryption of messages. Although the RFCs permit the use of either symmetric or asymmetric (public key) cryptography (for example, the RSA cryptosystem) to distribute symmetric keys, the RFCs strongly recommend the use of asymmetric cryptography for this purpose and to generate and validate digital signatures for messages and certificates. Public key management in PEM is based on the use of certificates as defined by the CCITT Directory Authentication Framework [CCIT88c]. A public key certification hierarchy for PEM is being established by the Internet Society. This certification hierarchy supports universal authentication of PEM users, under various policies, without the need for prior, bilateral arrangements among users or organizations with which the users may be affiliated.

Essay 18, "Electronic Data Interchange (EDI) Messaging Security," by Ted Humphreys, observes that modern economy and the future wealth and prosperity of industry and commerce rely increasingly on the exchange of data and information, in electronic form, between business partners. In response to the need for effective and efficient solutions to handle this way of doing business, Electronic Data Interchange (EDI) of-

¹The meaning of RFC has evolved. Today, most RFCs are effectively standards. Draft RFCs are used to solicit comments.

fers substantial advantages and opportunities. This essay looks at a particularly important aspect of EDI: the security of EDI messages. In particular, it focuses on the secure communications of EDI messages. To start with, some introductory material is presented that views security in the context of Open-EDI.

Essay 19, “Architectures for MLS Database Management Systems,” by LouAnna Notargiacomo, presents an overview of the basic architectures that have been used in the development of trusted relational database management systems (DBMSs). While various approaches have been tried for special-purpose systems, the architectures presented are those that have been developed for general-purpose trusted DBMS products. The essay also reviews approaches that have been proposed in the research for new trusted DBMS architectures, although worked examples of these approaches may not exist in all cases. Each component of the architecture is defined and the relationships and flow of information among components presented. This presentation is followed by a discussion of how the architecture meets mandatory and discretionary security requirements and preserves data integrity.

Essay 20, “Toward a Multilevel Secure Relational Data Model,” by Sushil Jajodia and Ravi Sandhu, observes that although there are several efforts under way to build multilevel secure relational database management systems, there is no clear consensus regarding what a multilevel secure relational data model exactly is. In part, this lack of consensus on fundamental issues reflects the subtleties involved in extending the classical (single-level) relational model to a multilevel environment. The authors’ aim in this essay is to discuss the most fundamental aspects of the multilevel secure relational model. First, they identify four core integrity properties that should be required of all multilevel relations. Next, they give a formal operational semantics for the usual update operations (*insert*, *update*, and *delete*) on multilevel relations. Finally, they describe a decomposition algorithm that partitions the multilevel relations into collections of single-level relations, and a recovery algorithm that constructs the original multilevel relations from the decomposed single-level relations.

Essay 21, “Solutions to the Polyinstantiation Problem,” by Sushil Jajodia, Ravi Sandhu, and Barbara Blaustein, addresses polyinstantiation, which has generated a great deal of controversy lately. Some have argued that polyinstantiation and integrity are fundamentally incompatible, and have proposed alternatives to polyinstantiation. Others have argued about the correct definition of polyinstantiation and its operational semantics. The purpose of this essay is to provide a tutorial on the subject. The authors begin by reviewing the concept of polyinstantiation; then they survey the various proposals to deal with it.

Essay 22, “Integrity in Multilevel Secure Database Management Systems,” by Catherine Meadows and Sushil Jajodia, discusses the effects

that satisfying security requirements in a multilevel database management system can have on the system's data integrity. The authors identify the conflicts between security and integrity in such databases, and show how the various components of integrity can be traded off both against each other and against security. They discuss recent work in maintaining integrity in multilevel relational database management systems and identify the emerging integrity issues in multilevel object-oriented systems.

In Essay 23, "Multilevel Secure Database Management Prototypes," Thomas Hinke describes, compares, and contrasts three of the most prominent research DBMS prototypes: SRI International's SeaView, TRW's Advanced Secure DBMS (ASD), and SCTC's Lock Data View (LDV). While each of these systems targets the A1 level of evaluation, they differ in the nature of the security policy enforced or in the architectural approach used to achieve their security. These systems represent a range of architectural approaches, with ASD taking a trusted process approach, and SeaView and LDV relying on the underlying trusted computing base (TCB) for mandatory security enforcement. These latter two systems thus provide an interesting contrast. LDV is also interesting because it enforces policies beyond those enforced in the other two systems.

In Essay 24, "Inference Problems in Multilevel Secure Database Management Systems," Sushil Jajodia and Catherine Meadows observe that inference is the process of deriving new information from known information. In multilevel database systems, the inference problem refers to the fact that the derived information can have higher sensitivity than the information provided to the user by the system. This essay surveys the state of the art in the study of inference problems. It defines and characterizes the inference problem as it relates to multilevel database systems and describes methods that have been developed for dealing with it.

Essay 25, "Logical Design of Audit Information in Relational Databases," by Sushil Jajodia, Shashi Gadia, and Gautam Bhargava, considers situations where the data is sufficiently sensitive that an audit trail becomes a necessity. Unfortunately, existing databases make a distinction between the current and past data. While they provide various types of support for dealing with the current data, the support for audit data is either nonexistent or very rudimentary. In this essay, the authors describe the database activity model that imposes a uniform logical structure on the past, present, and future data. There is never any loss of historical or current information in this model; thus the model provides a convenient mechanism for complete reconstruction of every action taken on the database.

Essay 26, "A Multilevel-Secure Object-Oriented Data Model," by Sushil Jajodia, Boris Kogan, and Ravi Sandhu, presents a new security model

for mandatory access controls in object-oriented database systems. This model is a departure from the traditional security models based on the passive object, active subject paradigm. The authors' model is a flow model whose main elements are objects and messages. An object combines the properties of a passive information repository with those of an active agent. Messages and their replies are the basic instruments of information flow. The chief advantages of the proposed model are its compatibility with the object-oriented data model and the simplicity with which mandatory security policies can be stated and enforced.

The goal of Essay 27, "Integrity Mechanisms in Database Management Systems," by Ravi Sandhu and Sushil Jajodia, is to answer the following question: What mechanisms are required in a general-purpose multi-user database management system (DBMS) to facilitate the integrity objectives of information systems? Although existing commercial products fall short of providing the requisite mechanisms, in principle they can be easily extended to incorporate these mechanisms.

In addition to the essays, we provide an extensive glossary and an index. Biographies of the editors and authors follow.