

DoD 5200.28-STD
Supersedes
CSC-STD-001-83, dtd 15 Aug 83
Library No. S225,711

DEPARTMENT OF DEFENSE STANDARD

DEPARTMENT OF
DEFENSE
TRUSTED COMPUTER
SYSTEM EVALUATION
CRITERIA

DECEMBER 1985

December 26, 1985

FOREWORD

This publication, DoD 5200.28-STD, "Department of Defense Trusted Computer System Evaluation Criteria," is issued under the authority of an in accordance with DoD Directive 5200.28, "Security Requirements for Automatic Data Processing (ADP) Systems," and in furtherance of responsibilities assigned by DoD Directive 5215.1, "Computer Security Evaluation Center." Its purpose is to provide technical hardware/firmware/software security criteria and associated technical evaluation methodologies in support of the overall ADP system security policy, evaluation and approval/accreditation responsibilities promulgated by DoD Directive 5200.28.

The provisions of this document apply to the Office of the Secretary of Defense (ASD), the Military Departments, the Organization of the Joint Chiefs of Staff, the Unified and Specified Commands, the Defense Agencies and activities administratively supported by OSD (hereafter called "DoD Components").

This publication is effective immediately and is mandatory for use by all DoD Components in carrying out ADP system technical security evaluation activities applicable to the processing and storage of classified and other sensitive DoD information and applications as set forth herein.

Recommendations for revisions to this publication are encouraged and will be reviewed biannually by the National Computer Security Center through a formal review process. Address all proposals for revision through appropriate channels to: National Computer Security Center, Attention: Chief, Computer Security Standards.

DoD Components may obtain copies of this publication through their own publications channels. Other federal agencies and the public may obtain copies from: Office of Standards and Products, National Computer Security Center, Fort Meade, MD 20755-6000, Attention: Chief, Computer Security Standards.

Donald C. Latham
Assistant Secretary of Defense
(Command, Control, Communications, and Intelligence)

ACKNOWLEDGEMENTS

Special recognition is extended to Sheila L. Brand, National Computer Security Center (NCSC), who integrated theory, policy, and practice into and directed the production of this document.

Acknowledgment is also given for the contributions of: Grace Hammonds and Peter S. Tasker, the MITRE Corp., Daniel J. Edwards, NCSC, Roger R. Schell, former Deputy Director of NCSC, Marvin Schaefer, NCSC, and Theodore M. P. Lee, Sperry Corp., who as original architects formulated and articulated the technical issues and solutions presented in this document; Jeff Makey, formerly NCSC, Warren F. Shadle, NCSC, and Carole S. Jordan, NCSC, who assisted in the preparation of this document; James P. Anderson, James P. Anderson & Co., Steven B. Lipner, Digital Equipment Corp., Clark Weissman, System Development Corp., LTC Lawrence A. Noble, formerly U.S. Air Force, Stephen T. Walker, formerly DoD, Eugene V. Epperly, DoD, and James E. Studer, formerly Dept. of the Army, who gave generously of their time and expertise in the review and critique of this document; and finally, thanks are given to the computer industry and others interested in trusted computing for their enthusiastic advice and assistance throughout this effort.

CONTENTS

FOREWORD.	.i
ACKNOWLEDGMENTS	ii
PREFACE	v
INTRODUCTION.	1

PART I: THE CRITERIA

1.0 DIVISION D: MINIMAL PROTECTION.	9
2.0 DIVISION C: DISCRETIONARY PROTECTION.	11
2.1 Class (C1): Discretionary Security Protection	12
2.2 Class (C2): Controlled Access Protection.	15
3.0 DIVISION B: MANDATORY PROTECTION.	19
3.1 Class (B1): Labeled Security Protection	20
3.2 Class (B2): Structured Protection	26
3.3 Class (B3): Security Domains.	33
4.0 DIVISION A: VERIFIED PROTECTION	41
4.1 Class (A1): Verified Design	42
4.2 Beyond Class (A1).	51

PART II: RATIONALE AND GUIDELINES

5.0 CONTROL OBJECTIVES FOR TRUSTED COMPUTER SYSTEMS.	55
5.1 A Need for Consensus	56
5.2 Definition and Usefulness.	56
5.3 Criteria Control Objective	56
6.0 RATIONALE BEHIND THE EVALUATION CLASSES.	63
6.1 The Reference Monitor Concept.	64
6.2 A Formal Security Policy Model	64
6.3 The Trusted Computing Base	65
6.4 Assurance.	65
6.5 The Classes.	66
7.0 THE RELATIONSHIP BETWEEN POLICY AND THE CRITERIA	69
7.1 Established Federal Policies	70
7.2 DoD Policies	70
7.3 Criteria Control Objective For Security Policy	71
7.4 Criteria Control Objective for Accountability.	74
7.5 Criteria Control Objective for Assurance	76
8.0 A GUIDELINE ON COVERT CHANNELS	79

9.0	A GUIDELINE ON CONFIGURING MANDATORY ACCESS CONTROL FEATURES	81
10.0	A GUIDELINE ON SECURITY TESTING	83
10.1	Testing for Division C	84
10.2	Testing for Division B	84
10.3	Testing for Division A	85
APPENDIX A: Commercial Product Evaluation Process.		87
APPENDIX B: Summary of Evaluation Criteria Divisions		89
APPENDIX C: Sumary of Evaluation Criteria Classes.		91
APPENDIX D: Requirement Directory.		93
GLOSSARY.		109
REFERENCES.		115

PREFACE

The trusted computer system evaluation criteria defined in this document classify systems into four broad hierarchical divisions of enhanced security protection. They provide a basis for the evaluation of effectiveness of security controls built into automatic data processing system products. The criteria were developed with three objectives in mind: (a) to provide users with a yardstick with which to assess the degree of trust that can be placed in computer systems for the secure processing of classified or other sensitive information; (b) to provide guidance to manufacturers as to what to build into their new, widely-available trusted commercial products in order to satisfy trust requirements for sensitive applications; and (c) to provide a basis for specifying security requirements in acquisition specifications. Two types of requirements are delineated for secure processing: (a) specific security feature requirements and (b) assurance requirements. Some of the latter requirements enable evaluation personnel to determine if the required features are present and functioning as intended. The scope of these criteria is to be applied to the set of components comprising a trusted system, and is not necessarily to be applied to each system component individually. Hence, some components of a system may be completely untrusted, while others may be individually evaluated to a lower or higher evaluation class than the trusted product considered as a whole system. In trusted products at the high end of the range, the strength of the reference monitor is such that most of the components can be completely untrusted. Though the criteria are intended to be application-independent, the specific security feature requirements may have to be interpreted when applying the criteria to specific systems with their own functional requirements, applications or special environments (e.g., communications processors, process control computers, and embedded systems in general). The underlying assurance requirements can be applied across the entire spectrum of ADP system or application processing environments without special interpretation.

INTRODUCTION

Historical Perspective

In October 1967, a task force was assembled under the auspices of the Defense Science Board to address computer security safeguards that would protect classified information in remote-access, resource-sharing computer systems. The Task Force report, "Security Controls for Computer Systems," published in February 1970, made a number of policy and technical recommendations on actions to be taken to reduce the threat of compromise of classified information processed on remote-access computer systems.[34] Department of Defense Directive 5200.28 and its accompanying manual DoD 5200.28-M, published in 1972 and 1973 respectively, responded to one of these recommendations by establishing uniform DoD policy, security requirements, administrative controls, and technical measures to protect classified information processed by DoD computer systems.[8;9] Research and development work undertaken by the Air Force, Advanced Research Projects Agency, and other defense agencies in the early and mid 70's developed and demonstrated solution approaches for the technical problems associated with controlling the flow of information in resource and information sharing computer systems.[1] The DoD Computer Security Initiative was started in 1977 under the auspices of the Under Secretary of Defense for Research and Engineering to focus DoD efforts addressing computer security issues.[33]

Concurrent with DoD efforts to address computer security issues, work was begun under the leadership of the National Bureau of Standards (NBS) to define problems and solutions for building, evaluating, and auditing secure computer systems.[17] As part of this work NBS held two invitational workshops on the subject of audit and evaluation of computer security.[20;28] The first was held in March 1977, and the second in November of 1978. One of the products of the second workshop was a definitive paper on the problems related to providing criteria for the evaluation of technical computer security effectiveness.[20] As an outgrowth of recommendations from this report, and in support of the DoD Computer Security Initiative, the MITRE Corporation began work on a set of computer security evaluation criteria that could be used to assess the degree of trust one could place in a computer system to protect classified data.[24;25;31] The preliminary concepts for computer security evaluation were defined and expanded upon at invitational workshops and symposia whose participants represented computer security expertise drawn from industry and academia in addition to the government. Their work has since been subjected to much peer review and constructive technical criticism from the DoD, industrial research and development organizations, universities, and computer manufacturers.

The DoD Computer Security Center (the Center) was formed in January 1981 to staff and expand on the work started by the DoD Computer Security Initiative.[15] A major goal of the Center as given in its DoD Charter is to encourage the widespread availability of trusted computer systems for use by those who process classified or other sensitive information.[10] The criteria presented in this document have evolved from the earlier NBS and MITRE evaluation material.

Scope

The trusted computer system evaluation criteria defined in this document apply primarily to trusted commercially available automatic data processing (ADP) systems. They are also applicable, as amplified below, to the evaluation of existing systems and to the specification of security requirements for ADP systems acquisition. Included are two distinct sets of requirements: 1) specific security feature requirements; and 2) assurance requirements. The specific feature requirements encompass the capabilities typically found in information processing systems employing general-purpose operating systems that are distinct from the applications programs being supported. However, specific security feature requirements may also apply to specific systems with their own functional requirements, applications or special environments (e.g., communications processors, process control computers, and embedded systems in general). The assurance requirements, on the other hand, apply to systems that cover the full range of computing environments from dedicated controllers to full range multilevel secure resource sharing systems.

Purpose

As outlined in the Preface, the criteria have been developed to serve a number of intended purposes:

- * To provide a standard to manufacturers as to what security features to build into their new and planned, commercial products in order to provide widely available systems that satisfy trust requirements (with particular emphasis on preventing the disclosure of data) for sensitive applications.
- * To provide DoD components with a metric with which to evaluate the degree of trust that can be placed in computer systems for the secure processing of classified and other sensitive information.
- * To provide a basis for specifying security requirements in acquisition specifications.

With respect to the second purpose for development of the criteria, i.e., providing DoD components with a security evaluation metric, evaluations can be delineated into two types: (a) an evaluation can be performed on a computer product from a perspective that excludes the application environment; or, (b) it can be done to assess whether appropriate security measures have been taken to permit the system to be used operationally in a specific environment. The former type of evaluation is done by the Computer Security Center through the Commercial Product Evaluation Process. That process is described in Appendix A.

The latter type of evaluation, i.e., those done for the purpose of assessing a system's security attributes with respect to a specific operational mission, is known as a certification evaluation. It must be understood that the completion of a formal product evaluation does not constitute certification or accreditation for the system to be used in any specific application environment. On the contrary, the evaluation report only provides a trusted computer system's evaluation rating along with supporting data describing the product system's strengths and weaknesses from a computer security point of

view. The system security certification and the formal approval/accreditation procedure, done in accordance with the applicable policies of the issuing agencies, must still be followed-before a system can be approved for use in processing or handling classified information.[8;9] Designated Approving Authorities (DAAs) remain ultimately responsible for specifying security of systems they accredit.

The trusted computer system evaluation criteria will be used directly and indirectly in the certification process. Along with applicable policy, it will be used directly as technical guidance for evaluation of the total system and for specifying system security and certification requirements for new acquisitions. Where a system being evaluated for certification employs a product that has undergone a Commercial Product Evaluation, reports from that process will be used as input to the certification evaluation. Technical data will be furnished to designers, evaluators and the Designated Approving Authorities to support their needs for making decisions.

Fundamental Computer Security Requirements

Any discussion of computer security necessarily starts from a statement of requirements, i.e., what it really means to call a computer system "secure." In general, secure systems will control, through use of specific security features, access to information such that only properly authorized individuals, or processes operating on their behalf, will have access to read, write, create, or delete information. Six fundamental requirements are derived from this basic statement of objective: four deal with what needs to be provided to control access to information; and two deal with how one can obtain credible assurances that this is accomplished in a trusted computer system.

Policy

Requirement 1 - SECURITY POLICY - There must be an explicit and well-defined security policy enforced by the system. Given identified subjects and objects, there must be a set of rules that are used by the system to determine whether a given subject can be permitted to gain access to a specific object. Computer systems of interest must enforce a mandatory security policy that can effectively implement access rules for handling sensitive (e.g., classified) information.[7] These rules include requirements such as: No person lacking proper personnel security clearance shall obtain access to classified information. In addition, discretionary security controls are required to ensure that only selected users or groups of users may obtain access to data (e.g., based on a need-to-know).

Requirement 2 - MARKING - Access control labels must be associated with objects. In order to control access to information stored in a computer, according to the rules of a mandatory security policy, it must be possible to mark every object with a label that reliably identifies the object's sensitivity level (e.g., classification), and/or the modes of access accorded those subjects who may potentially access the object.

Accountability

Requirement 3 - IDENTIFICATION - Individual subjects must be identified. Each access to information must be mediated based on who is accessing the information and what classes of information they are authorized to deal with. This identification and authorization information must be securely maintained by the computer system and be associated with every active element that performs some security-relevant action in the system.

Requirement 4 - ACCOUNTABILITY - Audit information must be selectively kept and protected so that actions affecting security can be traced to the responsible party. A trusted system must be able to record the occurrences of security-relevant events in an audit log. The capability to select the audit events to be recorded is necessary to minimize the expense of auditing and to allow efficient analysis. Audit data must be protected from modification and unauthorized destruction to permit detection and after-the-fact investigations of security violations.

Assurance

Requirement 5 - ASSURANCE - The computer system must contain hardware/software mechanisms that can be independently evaluated to provide sufficient assurance that the system enforces requirements 1 through 4 above. In order to assure that the four requirements of Security Policy, Marking, Identification, and Accountability are enforced by a computer system, there must be some identified and unified collection of hardware and software controls that perform those functions. These mechanisms are typically embedded in the operating system and are designed to carry out the assigned tasks in a secure manner. The basis for trusting such system mechanisms in their operational setting must be clearly documented such that it is possible to independently examine the evidence to evaluate their sufficiency.

Requirement 6 - CONTINUOUS PROTECTION - The trusted mechanisms that enforce these basic requirements must be continuously protected against tampering and/or unauthorized changes. No computer system can be considered truly secure if the basic hardware and software mechanisms that enforce the security policy are themselves subject to unauthorized modification or subversion. The continuous protection requirement has direct implications throughout the computer system's life-cycle.

These fundamental requirements form the basis for the individual evaluation criteria applicable for each evaluation division and class. The interested reader is referred to Section 5 of this document, "Control Objectives for Trusted Computer Systems," for a more complete discussion and further amplification of these fundamental requirements as they apply to general-purpose information processing systems and to Section 7 for amplification of the relationship between Policy and these requirements.

Structure of the Document

The remainder of this document is divided into two parts, four appendices, and a glossary. Part I (Sections 1 through 4) presents the detailed criteria derived from the fundamental requirements described above and relevant to the

rationale and policy excerpts contained in Part II.

Part II (Sections 5 through 10) provides a discussion of basic objectives, rationale, and national policy behind the development of the criteria, and guidelines for developers pertaining to: mandatory access control rules implementation, the covert channel problem, and security testing. It is divided into six sections. Section 5 discusses the use of control objectives in general and presents the three basic control objectives of the criteria. Section 6 provides the theoretical basis behind the criteria. Section 7 gives excerpts from pertinent regulations, directives, OMB Circulars, and Executive Orders which provide the basis for many trust requirements for processing nationally sensitive and classified information with computer systems. Section 8 provides guidance to system developers on expectations in dealing with the covert channel problem. Section 9 provides guidelines dealing with mandatory security. Section 10 provides guidelines for security testing. There are four appendices, including a description of the Trusted Computer System Commercial Products Evaluation Process (Appendix A), summaries of the evaluation divisions (Appendix B) and classes (Appendix C), and finally a directory of requirements ordered alphabetically. In addition, there is a glossary.

Structure of the Criteria

The criteria are divided into four divisions: D, C, B, and A ordered in a hierarchical manner with the highest division (A) being reserved for systems providing the most comprehensive security. Each division represents a major improvement in the overall confidence one can place in the system for the protection of sensitive information. Within divisions C and B there are a number of subdivisions known as classes. The classes are also ordered in a hierarchical manner with systems representative of division C and lower classes of division B being characterized by the set of computer security mechanisms that they possess. Assurance of correct and complete design and implementation for these systems is gained mostly through testing of the security-relevant portions of the system. The security-relevant portions of a system are referred to throughout this document as the Trusted Computing Base (TCB). Systems representative of higher classes in division B and division A derive their security attributes more from their design and implementation structure. Increased assurance that the required features are operative, correct, and tamperproof under all circumstances is gained through progressively more rigorous analysis during the design process.

Within each class, four major sets of criteria are addressed. The first three represent features necessary to satisfy the broad control objectives of Security Policy, Accountability, and Assurance that are discussed in Part II, Section 5. The fourth set, Documentation, describes the type of written evidence in the form of user guides, manuals, and the test and design documentation required for each class.

A reader using this publication for the first time may find it helpful to first read Part II, before continuing on with Part I.

PART I: THE CRITERIA

Highlighting (UPPERCASE) is used in Part I to indicate criteria not contained in a lower class or changes and additions to already defined criteria. Where there is no highlighting, requirements have been carried over from lower classes without addition or modification.

1.0 DIVISION D: MINIMAL PROTECTION

This division contains only one class. It is reserved for those systems that have been evaluated but that fail to meet the requirements for a higher evaluation class.

2.0 DIVISION C: DISCRETIONARY PROTECTION

Classes in this division provide for discretionary (need-to-know) protection and, through the inclusion of audit capabilities, for accountability of subjects and the actions they initiate.

2.1 CLASS (C1): DISCRETIONARY SECURITY PROTECTION

The Trusted Computing Base (TCB) of a class (C1) system nominally satisfies the discretionary security requirements by providing separation of users and data. It incorporates some form of credible controls capable of enforcing access limitations on an individual basis, i.e., ostensibly suitable for allowing users to be able to protect project or private information and to keep other users from accidentally reading or destroying their data. The class (C1) environment is expected to be one of cooperating users processing data at the same level(s) of sensitivity. The following are minimal requirements for systems assigned a class (C1) rating:

2.1.1 Security Policy

2.1.1.1 Discretionary Access Control

The TCB shall define and control access between named users and named objects (e.g., files and programs) in the ADP system. The enforcement mechanism (e.g., self/group/public controls, access control lists) shall allow users to specify and control sharing of those objects by named individuals or defined groups or both.

2.1.2 Accountability

2.1.2.1 Identification and Authentication

The TCB shall require users to identify themselves to it before beginning to perform any other actions that the TCB is expected to mediate. Furthermore, the TCB shall use a protected mechanism (e.g., passwords) to authenticate the user's identity. The TCB shall protect authentication data so that it cannot be accessed by any unauthorized user.

2.1.3 Assurance

2.1.3.1 Operational Assurance

2.1.3.1.1 System Architecture

The TCB shall maintain a domain for its own execution protects it from external interference or tampering (e.g., by modification of its code or data structures). Resources controlled by the TCB may be a defined subset of the subjects and objects in the ADP system.

2.1.3.1.2 System Integrity

Hardware and/or software features shall be provided that can be used to periodically validate the correct operation of the on-site hardware and firmware elements of the TCB.

2.1.3.2 Life-Cycle Assurance

2.1.3.2.1 Security Testing

The security mechanisms of the ADP system shall be tested and found to work as claimed in the system documentation. Testing shall be done to assure that there are no obvious ways for an unauthorized user to bypass or otherwise defeat the security protection mechanisms of the TCB. (See the Security Testing Guidelines.)

2.1.4 Documentation

2.1.4.1 Security Features User's Guide

A single summary, chapter, or manual in user documentation shall describe the protection mechanisms provided by the TCB, guidelines on their use, and how they interact with one another.

2.1.4.2 Trusted Facility Manual

A manual addressed to the ADP System Administrator shall present cautions about functions and privileges that should be controlled when running a secure facility.

2.1.4.3 Test Documentation

The system developer shall provide to the evaluators a document that describes the test plan, test procedures that show how the the security mechanisms were tested, and results of the security mechanisms' functional testing.

2.1.4.4 Design Documentation

Documentation shall be available that provides a description of the manufacturer's philosophy of protection and an explanation of how this philosophy is translated into the TCB. If the TCB is composed of distinct modules, the interfaces between these modules shall be described.

2.2 CLASS (C2): CONTROLLED ACCESS PROTECTION

Systems in this class enforce a more finely grained discretionary access control than (C1) systems, making users individually accountable for their actions through login procedures, auditing of security-relevant events, and resource isolation. The following are minimal requirements for systems assigned a class (C2) rating:

2.2.1 Security Policy

2.2.1.1 Discretionary Access Control

The TCB shall define and control access between named users and named objects (e.g., files and programs) in the ADP system. The enforcement mechanism (e.g., self/group/public controls, access control lists) shall allow users to specify and control sharing of those objects by named individuals, or defined groups of individuals, or by both, and shall provide controls to limit propagation of access rights. The discretionary access control mechanism shall, either by explicit user action or by default, provide that objects are protected from unauthorized access. These access controls shall be capable of including or excluding access to the granularity of a single user. Access permission to an object by users not already possessing access permission shall only be assigned by authorized users.

2.2.1.2 Object Reuse

All authorizations to the information contained within a storage object shall be revoked prior to initial assignment, allocation or reallocation to a subject from the TCB's pool of unused storage objects. No information, including encrypted representations of information, produced by a prior subject's actions is to be available to any subject that obtains access to an object that has been released back to the system.

2.2.2 Accountability

2.2.2.1 Identification and Authentication

The TCB shall require users to identify themselves to it before beginning to perform any other actions that the TCB is expected to mediate. Furthermore, the TCB shall use a protected mechanism (e.g., passwords) to authenticate the user's identity.

The TCB shall protect authentication data so that it cannot be accessed by any unauthorized user. The TCB shall be able to enforce individual accountability by providing the capability to uniquely identify each individual ADP system user. The TCB shall also provide the capability of associating this identity with all auditable actions taken by that individual.

2.2.2.2 Audit

The TCB shall be able to create, maintain, and protect from modification or unauthorized access or destruction an audit trail of accesses to the objects it protects. The audit data shall be protected by the TCB so that read access to it is limited to those who are authorized for audit data. The TCB shall be able to record the following types of events: use of identification and authentication mechanisms, introduction or objects into a user's address space (e.g., file open, program initiation), deletion of objects, and actions taken by computer operators and system administrators and/or system security officers, and other security relevant events. For each recorded event, the audit record shall identify: date and time of the event, user, type of event, and success or failure of the event. For identification/authentication events the origin of request (e.g., terminal ID) shall be included in the audit record. For events that introduce an object into a user's address space and for object deletion events the audit record shall include the name of the object. The ADP system administrator shall be able to selectively audit the actions of any one or more users based on individual identity.

2.2.3 Assurance

2.2.3.1 Operational Assurance

2.2.3.1.1 System Architecture

The TCB shall maintain a domain for its own execution that protects it from external interference or tampering (e.g., by modification of its code or data structures). Resources controlled by the TCB may be a defined subset of the subjects and objects in the ADP system. The TCB shall isolate the resources to be protected so that they are subject to the access control and auditing requirements.

2.2.3.1.2 System Integrity

Hardware and/or software features shall be provided that can be used to periodically validate the correct operation of the on-site hardware and firmware elements of the TCB.

2.2.3.2 Life-Cycle Assurance

2.2.3.2.1 Security Testing

The security mechanisms of the ADP system shall be tested and found to work as claimed in the system documentation. Testing shall be done to assure that there are no obvious ways for an unauthorized user to bypass or otherwise defeat the security protection mechanisms of the TCB. Testing shall also include a search for obvious flaws

that would allow violation of resource isolation, or that would permit unauthorized access to the audit or authentication data. (See the Security Testing guidelines.)

2.2.4 Documentation

2.2.4.1 Security Features User's Guide

A single summary, chapter, or manual in user documentation shall describe the protection mechanisms provided by the TCB, guidelines on their use, and how they interact with one another.

2.2.4.2 Trusted Facility Manual

A manual addressed to the ADP system administrator shall present cautions about functions and privileges that should be controlled when running a secure facility. The procedures for examining and maintaining the audit files as well as the detailed audit record structure for each type of audit event shall be given.

2.2.4.3 Test Documentation

The system developer shall provide to the evaluators a document that describes the test plan, test procedures that show how the security mechanisms were tested, and results of the security mechanisms' functional testing.

2.2.4.4 Design Documentation

Documentation shall be available that provides a description of the manufacturer's philosophy of protection and an explanation of how this philosophy is translated into the TCB. If the TCB is composed of distinct modules, the interfaces between these modules shall be described.

3.0 DIVISION B: MANDATORY PROTECTION

The notion of a TCB that preserves the integrity of sensitivity labels and uses them to enforce a set of mandatory access control rules is a major requirement in this division. Systems in this division must carry the sensitivity labels with major data structures in the system. The system developer also provides the security policy model on which the TCB is based and furnishes a specification of the TCB. Evidence must be provided to demonstrate that the reference monitor concept has been implemented.

3.1 CLASS (B1): LABELED SECURITY PROTECTION

Class (B1) systems require all the features required for class (C2). In addition, an informal statement of the security policy model, data labeling, and mandatory access control over named subjects and objects must be present. The capability must exist for accurately labeling exported information. Any flaws identified by testing must be removed. The following are minimal requirements for systems assigned a class (B1) rating:

3.1.1 Security Policy

3.1.1.1 Discretionary Access Control

The TCB shall define and control access between named users and named objects (e.g., files and programs) in the ADP system. The enforcement mechanism (e.g., self/group/public controls, access control lists) shall allow users to specify and control sharing of those objects by named individuals, or defined groups of individuals, or by both, and shall provide controls to limit propagation of access rights. The discretionary access control mechanism shall, either by explicit user action or by default, provide that objects are protected from unauthorized access. These access controls shall be capable of including or excluding access to the granularity of a single user. Access permission to an object by users not already possessing access permission shall only be assigned by authorized users.

3.1.1.2 Object Reuse

All authorizations to the information contained within a storage object shall be revoked prior to initial assignment, allocation or reallocation to a subject from the TCB's pool of unused storage objects. No information, including encrypted representations of information, produced by a prior subject's actions is to be available to any subject that obtains access to an object that has been released back to the system.

3.1.1.3 Labels

Sensitivity labels associated with each subject and storage object under its control (e.g., process, file, segment, device) shall be maintained by the TCB. These labels shall be used as the basis for mandatory access control decisions. In order to import non-labeled data, the TCB shall request and receive from an authorized user the security level of the data, and all such actions shall be auditable by the TCB.

3.1.1.3.1 Label Integrity

Sensitivity labels shall accurately represent security levels of the specific subjects or objects with which they are associated. When exported by the TCB,

sensitivity labels shall accurately and unambiguously represent the internal labels and shall be associated with the information being exported.

3.1.1.3.2 Exportation of Labeled Information

The TCB shall designate each communication channel and I/O device as either single-level or multilevel. Any change in this designation shall be done manually and shall be auditable by the TCB. The TCB shall maintain and be able to audit any change in the security level or levels associated with a communication channel or I/O device.

3.1.1.3.2.1 Exportation to Multilevel Devices

When the TCB exports an object to a multilevel I/O device, the sensitivity label associated with that object shall also be exported and shall reside on the same physical medium as the exported information and shall be in the same form (i.e., machine-readable or human-readable form). When the TCB exports or imports an object over a multilevel communication channel, the protocol used on that channel shall provide for the unambiguous pairing between the sensitivity labels and the associated information that is sent or received.

3.1.1.3.2.2 Exportation to Single-Level Devices

Single-level I/O devices and single-level communication channels are not required to maintain the sensitivity labels of the information they process. However, the TCB shall include a mechanism by which the TCB and an authorized user reliably communicate to designate the single security level of information imported or exported via single-level communication channels or I/O devices.

3.1.1.3.2.3 Labeling Human-Readable Output

The ADP system administrator shall be able to specify the printable label names associated with exported sensitivity labels. The TCB shall mark

* The hierarchical classification component in human-readable sensitivity labels shall be equal to the greatest hierarchical classification or any of the information in the output that the labels refer to; the non-hierarchical category component shall include all of the non-hierarchical categories of the information in the output the labels refer to, but no other non-hierarchical categories.

the beginning and end of all human-readable, paged, hardcopy output (e.g., line printer output) with human-readable sensitivity labels that properly* represent the sensitivity of the output. The TCB shall, by default, mark the top and bottom of each page of human-readable, paged, hardcopy output (e.g., line printer output) with human-readable sensitivity labels that properly* represent the overall sensitivity of the output or that properly* represent the sensitivity of the information on the page. The TCB shall, by default and in an appropriate manner, mark other forms of human-readable output (e.g., maps, graphics) with human-readable sensitivity labels that properly* represent the sensitivity of the output. Any override of these marking defaults shall be auditable by the TCB.

3.1.1.4 Mandatory Access Control

The TCB shall enforce a mandatory access control policy over all subjects and storage objects under its control (e.g., processes, files, segments, devices). These subjects and objects shall be assigned sensitivity labels that are a combination of hierarchical classification levels and non-hierarchical categories, and the labels shall be used as the basis for mandatory access control decisions. The TCB shall be able to support two or more such security levels. (See the Mandatory Access Control Guidelines.) The following requirements shall hold for all accesses between subjects and objects controlled by the TCB: a subject can read an object only if the hierarchical classification in the subject's security level is greater than or equal to the hierarchical classification in the object's security level and the non-hierarchical categories in the subject's security level include all the non-hierarchical categories in the object's security level. A subject can write an object only if the hierarchical classification in the subject's security level is less than or equal to the hierarchical classification in the object's security level and all the non-hierarchical categories in the subject's security level are included in the non-hierarchical categories in the object's security level. Identification and authentication data shall be used by the TCB to authenticate the user's identity and to ensure that the security level and authorization of subjects external to the TCB that may be created to act on behalf of the individual user are dominated by the clearance and authorization of that user.

3.1.2 Accountability

3.1.2.1 Identification and Authentication

The TCB shall require users to identify themselves to it before beginning to perform any other actions that the TCB is expected

to mediate. Furthermore, the TCB shall maintain authentication data that includes information for verifying the identity of individual users (e.g., passwords) as well as information for determining the clearance and authorizations of individual users. This data shall be used by the TCB to authenticate the user's identity and to ensure that the security level and authorizations of subjects external to the TCB that may be created to act on behalf of the individual user are dominated by the clearance and authorization of that user. The TCB shall protect authentication data so that it cannot be accessed by any unauthorized user. The TCB shall be able to enforce individual accountability by providing the capability to uniquely identify each individual ADP system user. The TCB shall also provide the capability of associating this identity with all auditable actions taken by that individual.

3.1.2.2 Audit

The TCB shall be able to create, maintain, and protect from modification or unauthorized access or destruction an audit trail of accesses to the objects it protects. The audit data shall be protected by the TCB so that read access to it is limited to those who are authorized for audit data. The TCB shall be able to record the following types of events: use of identification and authentication mechanisms, introduction of objects into a user's address space (e.g., file open, program initiation), deletion of objects, and actions taken by computer operators and system administrators and/or system security officers and other security relevant events. The TCB shall also be able to audit any override of human-readable output markings.

For each recorded event, the audit record shall identify: date and time of the event, user, type of event, and success or failure of the event. For identification/authentication events the origin of request (e.g., terminal ID) shall be included in the audit record. For events that introduce an object into a user's address space and for object deletion events the audit record shall include the name of the object and the object's security level. The ADP system administrator shall be able to selectively audit the actions of any one or more users based on individual identity and/or object security level.

3.1.3 Assurance

3.1.3.1 Operational Assurance

3.1.3.1.1 System Architecture

The TCB shall maintain a domain for its own execution that protects it from external interference or tampering (e.g., by modification of its code or data structures). Resources controlled by the TCB may be a defined subset of the subjects and objects in the ADP system. The TCB

shall maintain process isolation through the provision of distinct address spaces under its control. The TCB shall isolate the resources to be protected so that they are subject to the access control and auditing requirements.

3.1.3.1.2 System Integrity

Hardware and/or software features shall be provided that can be used to periodically validate the correct operation of the on-site hardware and firmware elements of the TCB.

3.1.3.2 Life-Cycle Assurance

3.1.3.2.1 Security Testing

The security mechanisms of the ADP system shall be tested and found to work as claimed in the system documentation. A team of individuals who thoroughly understand the specific implementation of the TCB shall subject its design documentation, source code, and object code to thorough analysis and testing. Their objectives shall be: to uncover all design and implementation flaws that would permit a subject external to the TCB to read, change, or delete data normally denied under the mandatory or discretionary security policy enforced by the TCB; as well as to assure that no subject (without authorization to do so) is able to cause the TCB to enter a state such that it is unable to respond to communications initiated by other users. All discovered flaws shall be removed or neutralized and the TCB retested to demonstrate that they have been eliminated and that new flaws have not been introduced. (See the Security Testing Guidelines.)

3.1.3.2.2 Design Specification and Verification

An informal or formal model of the security policy supported by the TCB shall be maintained over the life cycle of the ADP system and demonstrated to be consistent with its axioms.

3.1.4 Documentation

3.1.4.1 Security Features User's Guide

A single summary, chapter, or manual in user documentation shall describe the protection mechanisms provided by the TCB, guidelines on their use, and how they interact with one another.

3.1.4.2 Trusted Facility Manual

A manual addressed to the ADP system administrator shall

present cautions about functions and privileges that should be controlled when running a secure facility. The procedures for examining and maintaining the audit files as well as the detailed audit record structure for each type of audit event shall be given. The manual shall describe the operator and administrator functions related to security, to include changing the security characteristics of a user. It shall provide guidelines on the consistent and effective use of the protection features of the system, how they interact, how to securely generate a new TCB, and facility procedures, warnings, and privileges that need to be controlled in order to operate the facility in a secure manner.

3.1.4.3 Test Documentation

The system developer shall provide to the evaluators a document that describes the test plan, test procedures that show how the security mechanisms were tested, and results of the security mechanisms' functional testing.

3.1.4.4 Design Documentation

Documentation shall be available that provides a description of the manufacturer's philosophy of protection and an explanation of how this philosophy is translated into the TCB. If the TCB is composed of distinct modules, the interfaces between these modules shall be described. An informal or formal description of the security policy model enforced by the TCB shall be available and an explanation provided to show that it is sufficient to enforce the security policy. The specific TCB protection mechanisms shall be identified and an explanation given to show that they satisfy the model.

3.2 CLASS (B2): STRUCTURED PROTECTION

In class (B2) systems, the TCB is based on a clearly defined and documented formal security policy model that requires the discretionary and mandatory access control enforcement found in class (B1) systems be extended to all subjects and objects in the ADP system. In addition, covert channels are addressed. The TCB must be carefully structured into protection-critical and non- protection-critical elements. The TCB interface is well-defined and the TCB design and implementation enable it to be subjected to more thorough testing and more complete review. Authentication mechanisms are strengthened, trusted facility management is provided in the form of support for system administrator and operator functions, and stringent configuration management controls are imposed. The system is relatively resistant to penetration. The following are minimal requirements for systems assigned a class (B2) rating:

3.2.1 Security Policy

3.2.1.1 Discretionary Access Control

The TCB shall define and control access between named users and named objects (e.g., files and programs) in the ADP system. The enforcement mechanism (e.g., self/group/public controls, access control lists) shall allow users to specify and control sharing of those objects by named individuals, or defined groups of individuals, or by both, and shall provide controls to limit propagation of access rights. The discretionary access control mechanism shall, either by explicit user action or by default, provide that objects are protected from unauthorized access. These access controls shall be capable of including or excluding access to the granularity of a single user. Access permission to an object by users not already possessing access permission shall only be assigned by authorized users.

3.2.1.2 Object Reuse

All authorizations to the information contained within a storage object shall be revoked prior to initial assignment, allocation or reallocation to a subject from the TCB's pool of unused storage objects. No information, including encrypted representations of information, produced by a prior subject's actions is to be available to any subject that obtains access to an object that has been released back to the system.

3.2.1.3 Labels

Sensitivity labels associated with each ADP system resource (e.g., subject, storage object, ROM) that is directly or indirectly accessible by subjects external to the TCB shall be maintained by the TCB. These labels shall be used as the basis for mandatory access control decisions. In order to import non-labeled data, the TCB shall request and receive from an authorized user the security level of the data, and all such actions shall be auditable by the TCB.

3.2.1.3.1 Label Integrity

Sensitivity labels shall accurately represent security levels of the specific subjects or objects with which they are associated. When exported by the TCB, sensitivity labels shall accurately and unambiguously represent the internal labels and shall be associated with the information being exported.

3.2.1.3.2 Exportation of Labeled Information

The TCB shall designate each communication channel and I/O device as either single-level or multilevel. Any change in this designation shall be done manually and shall be auditable by the TCB. The TCB shall maintain and be able to audit any change in the security level or levels associated with a communication channel or I/O device.

3.2.1.3.2.1 Exportation to Multilevel Devices

When the TCB exports an object to a multilevel I/O device, the sensitivity label associated with that object shall also be exported and shall reside on the same physical medium as the exported information and shall be in the same form (i.e., machine-readable or human-readable form). When the TCB exports or imports an object over a multilevel communication channel, the protocol used on that channel shall provide for the unambiguous pairing between the sensitivity labels and the associated information that is sent or received.

3.2.1.3.2.2 Exportation to Single-Level Devices

Single-level I/O devices and single-level communication channels are not required to maintain the sensitivity labels of the information they process. However, the TCB shall include a mechanism by which the TCB and an authorized user reliably communicate to designate the single security level of information imported or exported via single-level communication channels or I/O devices.

3.2.1.3.2.3 Labeling Human-Readable Output

The ADP system administrator shall be able to specify the printable label names associated with exported sensitivity labels. The TCB shall mark the beginning and end of all human-readable, paged, hardcopy output (e.g., line printer output) with

human-readable sensitivity labels that properly* represent the sensitivity of the output. The TCB shall, by default, mark the top and bottom of each page of human-readable, paged, hardcopy output (e.g., line printer output) with human-readable sensitivity labels that properly* represent the overall sensitivity of the output or that properly* represent the sensitivity of the information on the page. The TCB shall, by default and in an appropriate manner, mark other forms of human-readable output (e.g., maps, graphics) with human-readable sensitivity labels that properly* represent the sensitivity of the output. Any override of these marking defaults shall be auditable by the TCB.

3.2.1.3.3 Subject Sensitivity Labels

The TCB shall immediately notify a terminal user of each change in the security level associated with that user during an interactive session. A terminal user shall be able to query the TCB as desired for a display of the subject's complete sensitivity label.

3.2.1.3.4 Device Labels

The TCB shall support the assignment of minimum and maximum security levels to all attached physical devices. These security levels shall be used by the TCB to enforce constraints imposed by the physical environments in which the devices are located.

3.2.1.4 Mandatory Access Control

The TCB shall enforce a mandatory access control policy over all resources (i.e., subjects, storage objects, and I/O devices that are directly or indirectly accessible by subjects external to the TCB. These subjects and objects shall be assigned sensitivity labels that are a combination of hierarchical classification levels and non-hierarchical categories, and the labels shall be used as the basis for mandatory access control decisions. The TCB shall be able to support two or more such security levels. (See the Mandatory Access Control guidelines.) The following requirements shall hold for all accesses between all subjects external to the TCB and all objects directly or indirectly accessible by these subjects: A subject can read an object only if the hierarchical classification in the subject's security level is greater than or equal to the hierarchical classification in the object's security level and the non-hierarchical categories in the subject's security level include all the non-hierarchical categories in the object's security level. A subject can write an object only if the hierarchical classification in the subject's security level is less than or

equal to the hierarchical classification in the object's security level and all the non-hierarchical categories in the subject's security level are included in the non-hierarchical categories in the object's security level. Identification and authentication data shall be used by the TCB to authenticate the user's identity and to ensure that the security level and authorization of subjects external to the TCB that may be created to act on behalf of the individual user are dominated by the clearance and authorization of that user.

3.2.2 Accountability

3.2.2.1 Identification and Authentication

The TCB shall require users to identify themselves to it before beginning to perform any other actions that the TCB is expected to mediate. Furthermore, the TCB shall maintain authentication data that includes information for verifying the identity of individual users (e.g., passwords) as well as information for determining the clearance and authorizations of individual users. This data shall be used by the TCB to authenticate the user's identity and to ensure that the security level and authorizations of subjects external to the TCB that may be created to act on behalf of the individual user are dominated by the clearance and authorization of that user. The TCB shall protect authentication data so that it cannot be accessed by any unauthorized user. The TCB shall be able to enforce individual accountability by providing the capability to uniquely identify each individual ADP system user. The TCB shall also provide the capability of associating this identity with all auditable actions taken by that individual.

3.2.2.1.1 Trusted Path

The TCB shall support a trusted communication path between itself and user for initial login and authentication. Communications via this path shall be initiated exclusively by a user.

3.2.2.2 Audit

The TCB shall be able to create, maintain, and protect from modification or unauthorized access or destruction an audit trail of accesses to the objects it protects. The audit data shall be protected by the TCB so that read access to it is limited to those who are authorized for audit data. The TCB shall be able to record the following types of events: use of identification and authentication mechanisms, introduction of objects into a user's address space (e.g., file open, program initiation), deletion of objects, and actions taken by computer operators and system administrators and/or system security officers, and other security relevant events. The TCB shall also be able to audit any override of human-readable output markings. For each recorded event, the audit record shall

identify: date and time of the event, user, type of event, and success or failure of the event. For identification/authentication events the origin of request (e.g., terminal ID) shall be included in the audit record. For events that introduce an object into a user's address space and for object deletion events the audit record shall include the name of the object and the object's security level. The ADP system administrator shall be able to selectively audit the actions of any one or more users based on individual identity and/or object security level. The TCB shall be able to audit the identified events that may be used in the exploitation of covert storage channels.

3.2.3 Assurance

3.2.3.1 Operational Assurance

3.2.3.1.1 System Architecture

The TCB shall maintain a domain for its own execution that protects it from external interference or tampering (e.g., by modification of its code or data structures). The TCB shall maintain process isolation through the provision of distinct address spaces under its control. The TCB shall be internally structured into well-defined largely independent modules. It shall make effective use of available hardware to separate those elements that are protection-critical from those that are not. The TCB modules shall be designed such that the principle of least privilege is enforced. Features in hardware, such as segmentation, shall be used to support logically distinct storage objects with separate attributes (namely: readable, writeable). The user interface to the TCB shall be completely defined and all elements of the TCB identified.

3.2.3.1.2 System Integrity

Hardware and/or software features shall be provided that can be used to periodically validate the correct operation of the on-site hardware and firmware elements of the TCB.

3.2.3.1.3 Covert Channel Analysis

The system developer shall conduct a thorough search for covert storage channels and make a determination (either by actual measurement or by engineering estimation) of the maximum bandwidth of each identified channel. (See the covert channels guideline section.)

3.2.3.1.4 Trusted Facility Management

The TCB shall support separate operator and administrator

functions.

3.2.3.2 Life-Cycle Assurance

3.2.3.2.1 Security Testing

The security mechanisms of the ADP system shall be tested and found to work as claimed in the system documentation. A team of individuals who thoroughly understand the specific implementation of the TCB shall subject its design documentation, source code, and object code to thorough analysis and testing. Their objectives shall be: to uncover all design and implementation flaws that would permit a subject external to the TCB to read, change, or delete data normally denied under the mandatory or discretionary security policy enforced by the TCB; as well as to assure that no subject (without authorization to do so) is able to cause the TCB to enter a state such that it is unable to respond to communications initiated by other users. The TCB shall be found relatively resistant to penetration. All discovered flaws shall be corrected and the TCB retested to demonstrate that they have been eliminated and that new flaws have not been introduced. Testing shall demonstrate that the TCB implementation is consistent with the descriptive top-level specification. (See the Security Testing Guidelines.)

3.2.3.2.2 Design Specification and Verification

A formal model of the security policy supported by the TCB shall be maintained over the life cycle of the ADP system that is proven consistent with its axioms. A descriptive top-level specification (DTLS) of the TCB shall be maintained that completely and accurately describes the TCB in terms of exceptions, error messages, and effects. It shall be shown to be an accurate description of the TCB interface.

3.2.3.2.3 Configuration Management

During development and maintenance of the TCB, a configuration management system shall be in place that maintains control of changes to the descriptive top-level specification, other design data, implementation documentation, source code, the running version of the object code, and test fixtures and documentation. The configuration management system shall assure a consistent mapping among all documentation and code associated with the current version of the TCB. Tools shall be provided for generation of a new version of the TCB from source code. Also available shall be tools for comparing a newly generated version with the previous TCB version in order to ascertain that only the intended changes have

been made in the code that will actually be used as the new version of the TCB.

3.2.4 Documentation

3.2.4.1 Security Features User's Guide

A single summary, chapter, or manual in user documentation shall describe the protection mechanisms provided by the TCB, guidelines on their use, and how they interact with one another.

3.2.4.2 Trusted Facility Manual

A manual addressed to the ADP system administrator shall present cautions about functions and privileges that should be controlled when running a secure facility. The procedures for examining and maintaining the audit files as well as the detailed audit record structure for each type of audit event shall be given. The manual shall describe the operator and administrator functions related to security, to include changing the security characteristics of a user. It shall provide guidelines on the consistent and effective use of the protection features of the system, how they interact, how to securely generate a new TCB, and facility procedures, warnings, and privileges that need to be controlled in order to operate the facility in a secure manner. The TCB modules that contain the reference validation mechanism shall be identified. The procedures for secure generation of a new TCB from source after modification of any modules in the TCB shall be described.

3.2.4.3 Test Documentation

The system developer shall provide to the evaluators a document that describes the test plan, test procedures that show how the security mechanisms were tested, and results of the security mechanisms' functional testing. It shall include results of testing the effectiveness of the methods used to reduce covert channel bandwidths.

3.2.4.4 Design Documentation

Documentation shall be available that provides a description of the manufacturer's philosophy of protection and an explanation of how this philosophy is translated into the TCB. The interfaces between the TCB modules shall be described. A formal description of the security policy model enforced by the TCB shall be available and proven that it is sufficient to enforce the security policy. The specific TCB protection mechanisms shall be identified and an explanation given to show that they satisfy the model. The descriptive top-level specification (DTLS) shall be shown to be an accurate description of the TCB interface. Documentation shall describe how the TCB implements the reference monitor concept and give

an explanation why it is tamper resistant, cannot be bypassed, and is correctly implemented. Documentation shall describe how the TCB is structured to facilitate testing and to enforce least privilege. This documentation shall also present the results of the covert channel analysis and the tradeoffs involved in restricting the channels. All auditable events that may be used in the exploitation of known covert storage channels shall be identified. The bandwidths of known covert storage channels the use of which is not detectable by the auditing mechanisms, shall be provided. (See the Covert Channel Guideline section.)

3.3 CLASS (B3): SECURITY DOMAINS

The class (B3) TCB must satisfy the reference monitor requirements that it mediate all accesses of subjects to objects, be tamperproof, and be small enough to be subjected to analysis and tests. To this end, the TCB is structured to exclude code not essential to security policy enforcement, with significant system engineering during TCB design and implementation directed toward minimizing its complexity. A security administrator is supported, audit mechanisms are expanded to signal security- relevant events, and system recovery procedures are required. The system is highly resistant to penetration. The following are minimal requirements for systems assigned a class (B3) rating:

3.1.1 Security Policy

3.3.1.1 Discretionary Access Control

The TCB shall define and control access between named users and named objects (e.g., files and programs) in the ADP system. The enforcement mechanism (e.g., access control lists) shall allow users to specify and control sharing of those objects, and shall provide controls to limit propagation of access rights. The discretionary access control mechanism shall, either by explicit user action or by default, provide that objects are protected from unauthorized access. These access controls shall be capable of specifying, for each named object, a list of named individuals and a list of groups of named individuals with their respective modes of access to that object. Furthermore, for each such named object, it shall be possible to specify a list of named individuals and a list of groups of named individuals for which no access to the object is to be given. Access permission to an object by users not already possessing access permission shall only be assigned by authorized users.

3.3.1.2 Object Reuse

All authorizations to the information contained within a storage object shall be revoked prior to initial assignment, allocation or reallocation to a subject from the TCB's pool of unused storage objects. No information, including encrypted representations of information, produced by a prior subjects actions is to be available to any subject that obtains access to an object that has been released back to the system.

3.3.1.3 Labels

Sensitivity labels associated with each ADP system resource (e.g., subject, storage object, ROM) that is directly or indirectly accessible by subjects external to the TCB shall be maintained by the TCB. These labels shall be used as the basis for mandatory access control decisions. In order to import non-labeled data, the TCB shall request and receive from an authorized user the security level of the data, and all such

hold for all accesses between all subjects external to the TCB and all objects directly or indirectly accessible by these actions shall be auditable by the TCB.

3.3.1.3.1 Label Integrity

Sensitivity labels shall accurately represent security levels of the specific subjects or objects with which they are associated. When exported by the TCB, sensitivity labels shall accurately and unambiguously represent the internal labels and shall be associated with the information being exported.

3.3.1.3.2 Exportation of Labeled Information

The TCB shall designate each communication channel and I/O device as either single-level or multilevel. Any change in this designation shall be done manually and shall be auditable by the TCB. The TCB shall maintain and be able to audit any change in the security level or levels associated with a communication channel or I/O device.

3.3.1.3.2.1 Exportation to Multilevel Devices

When the TCB exports an object to a multilevel I/O device, the sensitivity label associated with that object shall also be exported and shall reside on the same physical medium as the exported information and shall be in the same form (i.e., machine-readable or human-readable form). When the TCB exports or imports an object over a multilevel communication channel, the protocol used on that channel shall provide for the unambiguous pairing between the sensitivity labels and the associated information that is sent or received.

3.3.1.3.2.2 Exportation to Single-Level Devices

Single-level I/O devices and single-level communication channels are not required to maintain the sensitivity labels of the information they process. However, the TCB shall include a mechanism by which the TCB and an authorized user reliably communicate to designate the single security level of information imported or exported via single-level communication channels or I/O devices.

3.3.1.3.2.3 Labeling Human-Readable Output

The ADP system administrator shall be able to specify the printable label names associated with

exported sensitivity labels. The TCB shall mark the beginning and end of all human-readable, paged, hardcopy output (e.g., line printer output) with human-readable sensitivity labels that properly* represent the sensitivity of the output. The TCB shall, by default, mark the top and bottom of each page of human-readable, paged, hardcopy output (e.g., line printer output) with human-readable sensitivity labels that properly* represent the overall sensitivity of the output or that properly* represent the sensitivity of the information on the page. The TCB shall, by default and in an appropriate manner, mark other forms of human-readable output (e.g., maps, graphics) with human-readable sensitivity labels that properly* represent the sensitivity of the output. Any override of these marking defaults shall be auditable by the TCB.

3.3.1.3.3 Subject Sensitivity Labels

The TCB shall immediately notify a terminal user of each change in the security level associated with that user during an interactive session. A terminal user shall be able to query the TCB as desired for a display of the subject's complete sensitivity label.

3.3.1.3.4 Device Labels

The TCB shall support the assignment of minimum and maximum security levels to all attached physical devices. These security levels shall be used by the TCB to enforce constraints imposed by the physical environments in which the devices are located.

3.3.1.4 Mandatory Access Control

The TCB shall enforce a mandatory access control policy over all resources (i.e., subjects, storage objects, and I/O devices) that are directly or indirectly accessible by subjects external to the TCB. These subjects and objects shall be assigned sensitivity labels that are a combination of hierarchical classification levels and non-hierarchical categories, and the labels shall be used as the basis for mandatory access control decisions. The TCB shall be able to support two or more such security levels. (See the Mandatory

* The hierarchical classification component in human-readable sensitivity labels shall be equal to the greatest hierarchical classification of any of the information in the output that the labels refer to; the non-hierarchical category component shall include all of the non-hierarchical categories of the information in the output the labels refer to, but no other non-hierarchical categories.

Access Control guidelines.) The following requirements shall subjects: A subject can read an object only if the hierarchical classification in the subject's security level is greater than or equal to the hierarchical classification in the object's security level and the non-hierarchical categories in the subject's security level include all the non-hierarchical categories in the object's security level. A subject can write an object only if the hierarchical classification in the subject's security level is less than or equal to the hierarchical classification in the object's security level and all the non-hierarchical categories in the subject's security level are included in the non-hierarchical categories in the object's security level. Identification and authentication data shall be used by the TCB to authenticate the user's identity and to ensure that the security level and authorization of subjects external to the TCB that may be created to act on behalf of the individual user are dominated by the clearance and authorization of that user.

3.3.2 Accountability

3.3.2.1 Identification and Authentication

The TCB shall require users to identify themselves to it before beginning to perform any other actions that the TCB is expected to mediate. Furthermore, the TCB shall maintain authentication data that includes information for verifying the identity of individual users (e.g., passwords) as well as information for determining the clearance and authorizations of individual users. This data shall be used by the TCB to authenticate the user's identity and to ensure that the security level and authorizations of subjects external to the TCB that may be created to act on behalf of the individual user are dominated by the clearance and authorization of that user. The TCB shall protect authentication data so that it cannot be accessed by any unauthorized user. The TCB shall be able to enforce individual accountability by providing the capability to uniquely identify each individual ADP system user. The TCB shall also provide the capability of associating this identity with all auditable actions taken by that individual.

3.3.2.1.1 Trusted Path

The TCB shall support a trusted communication path between itself and users for use when a positive TCB-to-user connection is required (e.g., login, change subject security level). Communications via this trusted path shall be activated exclusively by a user of the TCB and shall be logically isolated and unmistakably distinguishable from other paths.

3.3.2.2 Audit

The TCB shall be able to create, maintain, and protect from

modification or unauthorized access or destruction an audit trail of accesses to the objects it protects. The audit data shall be protected by the TCB so that read access to it is limited to those who are authorized for audit data. The TCB shall be able to record the following types of events: use of identification and authentication mechanisms, introduction of objects into a user's address space (e.g., file open, program initiation), deletion of objects, and actions taken by computer operators and system administrators and/or system security officers and other security relevant events. The TCB shall also be able to audit any override of human-readable output markings. For each recorded event, the audit record shall identify: date and time of the event, user, type of event, and success or failure of the event. For identification/authentication events the origin of request (e.g., terminal ID) shall be included in the audit record. For events that introduce an object into a user's address space and for object deletion events the audit record shall include the name of the object and the object's security level. The ADP system administrator shall be able to selectively audit the actions of any one or more users based on individual identity and/or object security level. The TCB shall be able to audit the identified events that may be used in the exploitation of covert storage channels. The TCB shall contain a mechanism that is able to monitor the occurrence or accumulation of security auditable events that may indicate an imminent violation of security policy. This mechanism shall be able to immediately notify the security administrator when thresholds are exceeded, and if the occurrence or accumulation of these security relevant events continues, the system shall take the least disruptive action to terminate the event.

3.3.3 Assurance

3.3.3.1 Operational Assurance

3.3.3.1.1 System Architecture

The TCB shall maintain a domain for its own execution that protects it from external interference or tampering (e.g., by modification of its code or data structures). The TCB shall maintain process isolation through the provision of distinct address spaces under its control. The TCB shall be internally structured into well-defined largely independent modules. It shall make effective use of available hardware to separate those elements that are protection-critical from those that are not. The TCB modules shall be designed such that the principle of least privilege is enforced. Features in hardware, such as segmentation, shall be used to support logically distinct storage objects with separate attributes (namely: readable, writeable). The user interface to the TCB shall be completely defined and all elements of the TCB identified. The TCB shall be designed and structured

to use a complete, conceptually simple protection mechanism with precisely defined semantics. This mechanism shall play a central role in enforcing the internal structuring of the TCB and the system. The TCB shall incorporate significant use of layering, abstraction and data hiding. Significant system engineering shall be directed toward minimizing the complexity of the TCB and excluding from the TCB modules that are not protection-critical.

3.3.3.1.2 System Integrity

Hardware and/or software features shall be provided that can be used to periodically validate the correct operation of the on-site hardware and firmware elements of the TCB.

3.3.3.1.3 Covert Channel Analysis

The system developer shall conduct a thorough search for covert channels and make a determination (either by actual measurement or by engineering estimation) of the maximum bandwidth of each identified channel. (See the Covert Channels Guideline section.)

3.3.3.1.4 Trusted Facility Management

The TCB shall support separate operator and administrator functions. The functions performed in the role of a security administrator shall be identified. The ADP system administrative personnel shall only be able to perform security administrator functions after taking a distinct auditable action to assume the security administrator role on the ADP system. Non-security functions that can be performed in the security administration role shall be limited strictly to those essential to performing the security role effectively.

3.3.3.1.5 Trusted Recovery

Procedures and/or mechanisms shall be provided to assure that, after an ADP system failure or other discontinuity, recovery without a protection compromise is obtained.

3.3.3.2 Life-Cycle Assurance

3.3.3.2.1 Security Testing

The security mechanisms of the ADP system shall be tested and found to work as claimed in the system documentation. A team of individuals who thoroughly understand the specific implementation of the TCB shall subject its design documentation, source code, and object code to thorough analysis and testing. Their objectives shall

be: to uncover all design and implementation flaws that would permit a subject external to the TCB to read, change, or delete data normally denied under the mandatory or discretionary security policy enforced by the TCB; as well as to assure that no subject (without authorization to do so) is able to cause the TCB to enter a state such that it is unable to respond to communications initiated by other users. The TCB shall be found resistant to penetration. All discovered flaws shall be corrected and the TCB retested to demonstrate that they have been eliminated and that new flaws have not been introduced. Testing shall demonstrate that the TCB implementation is consistent with the descriptive top-level specification. (See the Security Testing Guidelines.) No design flaws and no more than a few correctable implementation flaws may be found during testing and there shall be reasonable confidence that few remain.

3.3.3.2.2 Design Specification and Verification

A formal model of the security policy supported by the TCB shall be maintained over the life cycle of the ADP system that is proven consistent with its axioms. A descriptive top-level specification (DTLS) of the TCB shall be maintained that completely and accurately describes the TCB in terms of exceptions, error messages, and effects. It shall be shown to be an accurate description of the TCB interface. A convincing argument shall be given that the DTLS is consistent with the model.

3.3.3.2.3 Configuration Management

During development and maintenance of the TCB, a configuration management system shall be in place that maintains control of changes to the descriptive top-level specification, other design data, implementation documentation, source code, the running version of the object code, and test fixtures and documentation. The configuration management system shall assure a consistent mapping among all documentation and code associated with the current version of the TCB. Tools shall be provided for generation of a new version of the TCB from source code. Also available shall be tools for comparing a newly generated version with the previous TCB version in order to ascertain that only the intended changes have been made in the code that will actually be used as the new version of the TCB.

3.3.4 Documentation

3.3.4.1 Security Features User's Guide

A single summary, chapter, or manual in user documentation shall describe the protection mechanisms provided by the TCB, guidelines on their use, and how they interact with one another.

3.3.4.2 Trusted Facility Manual

A manual addressed to the ADP system administrator shall present cautions about functions and privileges that should be controlled when running a secure facility. The procedures for examining and maintaining the audit files as well as the detailed audit record structure for each type of audit event shall be given. The manual shall describe the operator and administrator functions related to security, to include changing the security characteristics of a user. It shall provide guidelines on the consistent and effective use of the protection features of the system, how they interact, how to securely generate a new TCB, and facility procedures, warnings, and privileges that need to be controlled in order to operate the facility in a secure manner. The TCB modules that contain the reference validation mechanism shall be identified. The procedures for secure generation of a new TCB from source after modification of any modules in the TCB shall be described. It shall include the procedures to ensure that the system is initially started in a secure manner. Procedures shall also be included to resume secure system operation after any lapse in system operation.

3.3.4.3 Test Documentation

The system developer shall provide to the evaluators a document that describes the test plan, test procedures that show how the security mechanisms were tested, and results of the security mechanisms' functional testing. It shall include results of testing the effectiveness of the methods used to reduce covert channel bandwidths.

3.3.4.4 Design Documentation

Documentation shall be available that provides a description of the manufacturer's philosophy of protection and an explanation of how this philosophy is translated into the TCB. The interfaces between the TCB modules shall be described. A formal description of the security policy model enforced by the TCB shall be available and proven that it is sufficient to enforce the security policy. The specific TCB protection mechanisms shall be identified and an explanation given to show that they satisfy the model. The descriptive top-level specification (DTLS) shall be shown to be an accurate description of the TCB interface. Documentation shall describe how the TCB implements the reference monitor concept and give an explanation why it is tamper resistant, cannot be bypassed, and is correctly implemented. The TCB implementation (i.e., in hardware, firmware, and software) shall be informally shown to

be consistent with the DTLS. The elements of the DTLS shall be shown, using informal techniques, to correspond to the elements of the TCB. Documentation shall describe how the TCB is structured to facilitate testing and to enforce least privilege. This documentation shall also present the results of the covert channel analysis and the tradeoffs involved in restricting the channels. All auditable events that may be used in the exploitation of known covert storage channels shall be identified. The bandwidths of known covert storage channels, the use of which is not detectable by the auditing mechanisms, shall be provided. (See the Covert Channel Guideline section.)

4.0 DIVISION A: VERIFIED PROTECTION

This division is characterized by the use of formal security verification methods to assure that the mandatory and discretionary security controls employed in the system can effectively protect classified or other sensitive information stored or processed by the system. Extensive documentation is required to demonstrate that the TCB meets the security requirements in all aspects of design, development and implementation.

4.1 CLASS (A1): VERIFIED DESIGN

Systems in class (A1) are functionally equivalent to those in class (B3) in that no additional architectural features or policy requirements are added. The distinguishing feature of systems in this class is the analysis derived from formal design specification and verification techniques and the resulting high degree of assurance that the TCB is correctly implemented. This assurance is developmental in nature, starting with a formal model of the security policy and a formal top-level specification (FTLS) of the design. Independent of the particular specification language or verification system used, there are five important criteria for class (A1) design verification:

- * A formal model of the security policy must be clearly identified and documented, including a mathematical proof that the model is consistent with its axioms and is sufficient to support the security policy.
- * An FTLS must be produced that includes abstract definitions of the functions the TCB performs and of the hardware and/or firmware mechanisms that are used to support separate execution domains.
- * The FTLS of the TCB must be shown to be consistent with the model by formal techniques where possible (i.e., where verification tools exist) and informal ones otherwise.
- * The TCB implementation (i.e., in hardware, firmware, and software) must be informally shown to be consistent with the FTLS. The elements of the FTLS must be shown, using informal techniques, to correspond to the elements of the TCB. The FTLS must express the unified protection mechanism required to satisfy the security policy, and it is the elements of this protection mechanism that are mapped to the elements of the TCB.
- * Formal analysis techniques must be used to identify and analyze covert channels. Informal techniques may be used to identify covert timing channels. The continued existence of identified covert channels in the system must be justified.

In keeping with the extensive design and development analysis of the TCB required of systems in class (A1), more stringent configuration management is required and procedures are established for securely distributing the system to sites. A system security administrator is supported.

The following are minimal requirements for systems assigned a class (A1) rating:

4.1.1 Security Policy

4.1.1.1 Discretionary Access Control

The TCB shall define and control access between named users and named objects (e.g., files and programs) in the ADP system.

The enforcement mechanism (e.g., access control lists) shall allow users to specify and control sharing of those objects, and shall provide controls to limit propagation of access rights. The discretionary access control mechanism shall, either by explicit user action or by default, provide that objects are protected from unauthorized access. These access controls shall be capable of specifying, for each named object, a list of named individuals and a list of groups of named individuals with their respective modes of access to that object. Furthermore, for each such named object, it shall be possible to specify a list of named individuals and a list of groups of named individuals for which no access to the object is to be given. Access permission to an object by users not already possessing access permission shall only be assigned by authorized users.

4.1.1.2 Object Reuse

All authorizations to the information contained within a storage object shall be revoked prior to initial assignment, allocation or reallocation to a subject from the TCB's pool of unused storage objects. No information, including encrypted representations of information, produced by a prior subject's actions is to be available to any subject that obtains access to an object that has been released back to the system.

4.1.1.3 Labels

Sensitivity labels associated with each ADP system resource (e.g., subject, storage object, ROM) that is directly or indirectly accessible by subjects external to the TCB shall be maintained by the TCB. These labels shall be used as the basis for mandatory access control decisions. In order to import non-labeled data, the TCB shall request and receive from an authorized user the security level of the data, and all such actions shall be auditable by the TCB.

4.1.1.3.1 Label Integrity

Sensitivity labels shall accurately represent security levels of the specific subjects or objects with which they are associated. When exported by the TCB, sensitivity labels shall accurately and unambiguously represent the internal labels and shall be associated with the information being exported.

4.1.1.3.2 Exportation of Labeled Information

The TCB shall designate each communication channel and I/O device as either single-level or multilevel. Any change in this designation shall be done manually and shall be auditable by the TCB. The TCB shall maintain and be able to audit any change in the security level or levels associated with a communication channel or

I/O device.

4.1.1.3.2.1 Exportation to Multilevel Devices

When the TCB exports an object to a multilevel I/O device, the sensitivity label associated with that object shall also be exported and shall reside on the same physical medium as the exported information and shall be in the same form (i.e., machine-readable or human-readable form). When the TCB exports or imports an object over a multilevel communication channel, the protocol used on that channel shall provide for the unambiguous pairing between the sensitivity labels and the associated information that is sent or received.

4.1.1.3.2.2 Exportation to Single-Level Devices

Single-level I/O devices and single-level communication channels are not required to maintain the sensitivity labels of the information they process. However, the TCB shall include a mechanism by which the TCB and an authorized user reliably communicate to designate the single security level of information imported or exported via single-level communication channels or I/O devices.

4.1.1.3.2.3 Labeling Human-Readable Output

The ADP system administrator shall be able to specify the printable label names associated with exported sensitivity labels. The TCB shall mark the beginning and end of all human-readable, paged, hardcopy output (e.g., line printer output) with human-readable sensitivity labels that properly* represent the sensitivity of the output. The TCB shall, by default, mark the top and bottom of each page of human-readable, paged, hardcopy output (e.g., line printer output) with human-readable sensitivity labels that properly* represent the overall sensitivity of the output or that properly* represent the sensitivity of the information on the page. The TCB shall, by default and in an appropriate

* The hierarchical classification component in human-readable sensitivity labels shall be equal to the greatest hierarchical classification of any of the information in the output that the labels refer to; the non-hierarchical category component shall include all of the non-hierarchical categories of the information in the output the labels refer to, but no other non-hierarchical categories.

manner, mark other forms of human-readable output (e.g., maps, graphics) with human-readable sensitivity labels that properly* represent the sensitivity of the output. Any override of these marking defaults shall be auditable by the TCB.

4.1.1.3.3 Subject Sensitivity Labels

The TCB shall immediately notify a terminal user of each change in the security level associated with that user during an interactive session. A terminal user shall be able to query the TCB as desired for a display of the subject's complete sensitivity label.

4.1.1.3.4 Device Labels

The TCB shall support the assignment of minimum and maximum security levels to all attached physical devices. These security levels shall be used by the TCB to enforce constraints imposed by the physical environments in which the devices are located.

4.1.1.4 Mandatory Access Control

The TCB shall enforce a mandatory access control policy over all resources (i.e., subjects, storage objects, and I/O devices) that are directly or indirectly accessible by subjects external to the TCB. These subjects and objects shall be assigned sensitivity labels that are a combination of hierarchical classification levels and non-hierarchical categories, and the labels shall be used as the basis for mandatory access control decisions. The TCB shall be able to support two or more such security levels. (See the Mandatory Access Control guidelines.) The following requirements shall hold for all accesses between all subjects external to the TCB and all objects directly or indirectly accessible by these subjects: A subject can read an object only if the hierarchical classification in the subject's security level is greater than or equal to the hierarchical classification in the object's security level and the non-hierarchical categories in the subject's security level include all the non-hierarchical categories in the object's security level. A subject can write an object only if the hierarchical classification in the subject's security level is less than or equal to the hierarchical classification in the object's security level and all the non-hierarchical categories in the subject's security level are included in the non-hierarchical categories in the object's security level. Identification and authentication data shall be used by the TCB to authenticate the user's identity and to ensure that the security level and authorization of subjects external to the TCB that may be created to act on behalf of the individual user are dominated by the clearance and authorization of that user.

4.1.2 Accountability

4.1.2.1 Identification and Authentication

The TCB shall require users to identify themselves to it before beginning to perform any other actions that the TCB is expected to mediate. Furthermore, the TCB shall maintain authentication data that includes information for verifying the identity of individual users (e.g., passwords) as well as information for determining the clearance and authorizations of individual users. This data shall be used by the TCB to authenticate the user's identity and to ensure that the security level and authorizations of subjects external to the TCB that may be created to act on behalf of the individual user are dominated by the clearance and authorization of that user. The TCB shall protect authentication data so that it cannot be accessed by any unauthorized user. The TCB shall be able to enforce individual accountability by providing the capability to uniquely identify each individual ADP system user. The TCB shall also provide the capability of associating this identity with all auditable actions taken by that individual.

4.1.2.1.1 Trusted Path

The TCB shall support a trusted communication path between itself and users for use when a positive TCB-to-user connection is required (e.g., login, change subject security level). Communications via this trusted path shall be activated exclusively by a user or the TCB and shall be logically isolated and unmistakably distinguishable from other paths.

4.1.2.2 Audit

The TCB shall be able to create, maintain, and protect from modification or unauthorized access or destruction an audit trail of accesses to the objects it protects. The audit data shall be protected by the TCB so that read access to it is limited to those who are authorized for audit data. The TCB shall be able to record the following types of events: use of identification and authentication mechanisms, introduction of objects into a user's address space (e.g., file open, program initiation), deletion of objects, and actions taken by computer operators and system administrators and/or system security officers, and other security relevant events. The TCB shall also be able to audit any override of human-readable output markings. For each recorded event, the audit record shall identify: date and time of the event, user, type of event, and success or failure of the event. For identification/authentication events the origin of request (e.g., terminal ID) shall be included in the audit record. For events that introduce an object into a user's address space and for object deletion events the audit record shall include the name of the

object and the object's security level. The ADP system administrator shall be able to selectively audit the actions of any one or more users based on individual identity and/or object security level. The TCB shall be able to audit the identified events that may be used in the exploitation of covert storage channels. The TCB shall contain a mechanism that is able to monitor the occurrence or accumulation of security auditable events that may indicate an imminent violation of security policy. This mechanism shall be able to immediately notify the security administrator when thresholds are exceeded, and, if the occurrence or accumulation of these security relevant events continues, the system shall take the least disruptive action to terminate the event.

4.1.3 Assurance

4.1.3.1 Operational Assurance

4.1.3.1.1 System Architecture

The TCB shall maintain a domain for its own execution that protects it from external interference or tampering (e.g., by modification of its code or data structures). The TCB shall maintain process isolation through the provision of distinct address spaces under its control. The TCB shall be internally structured into well-defined largely independent modules. It shall make effective use of available hardware to separate those elements that are protection-critical from those that are not. The TCB modules shall be designed such that the principle of least privilege is enforced. Features in hardware, such as segmentation, shall be used to support logically distinct storage objects with separate attributes (namely: readable, writeable). The user interface to the TCB shall be completely defined and all elements of the TCB identified. The TCB shall be designed and structured to use a complete, conceptually simple protection mechanism with precisely defined semantics. This mechanism shall play a central role in enforcing the internal structuring of the TCB and the system. The TCB shall incorporate significant use of layering, abstraction and data hiding. Significant system engineering shall be directed toward minimizing the complexity of the TCB and excluding from the TCB modules that are not protection-critical.

4.1.3.1.2 System Integrity

Hardware and/or software features shall be provided that can be used to periodically validate the correct operation of the on-site hardware and firmware elements of the TCB.

4.1.3.1.3 Covert Channel Analysis

The system developer shall conduct a thorough search for covert channels and make a determination (either by actual measurement or by engineering estimation) of the maximum bandwidth of each identified channel. (See the Covert Channels Guideline section.) Formal methods shall be used in the analysis.

4.1.3.1.4 Trusted Facility Management

The TCB shall support separate operator and administrator functions. The functions performed in the role of a security administrator shall be identified. The ADP system administrative personnel shall only be able to perform security administrator functions after taking a distinct auditable action to assume the security administrator role on the ADP system. Non-security functions that can be performed in the security administration role shall be limited strictly to those essential to performing the security role effectively.

4.1.3.1.5 Trusted Recovery

Procedures and/or mechanisms shall be provided to assure that, after an ADP system failure or other discontinuity, recovery without a protection compromise is obtained.

4.1.3.2 Life-Cycle Assurance

4.1.3.2.1 Security Testing

The security mechanisms of the ADP system shall be tested and found to work as claimed in the system documentation. A team of individuals who thoroughly understand the specific implementation of the TCB shall subject its design documentation, source code, and object code to thorough analysis and testing. Their objectives shall be: to uncover all design and implementation flaws that would permit a subject external to the TCB to read, change, or delete data normally denied under the mandatory or discretionary security policy enforced by the TCB; as well as to assure that no subject (without authorization to do so) is able to cause the TCB to enter a state such that it is unable to respond to communications initiated by other users. The TCB shall be found resistant to penetration. All discovered flaws shall be corrected and the TCB retested to demonstrate that they have been eliminated and that new flaws have not been introduced. Testing shall demonstrate that the TCB implementation is consistent with the formal top-level specification. (See the Security Testing Guidelines.) No design flaws and no more than a few correctable implementation flaws may be found during testing and there shall be reasonable confidence that few

remain. Manual or other mapping of the FTLS to the source code may form a basis for penetration testing.

4.1.3.2.2 Design Specification and Verification

A formal model of the security policy supported by the TCB shall be maintained over the life-cycle of the ADP system that is proven consistent with its axioms. A descriptive top-level specification (DTLS) of the TCB shall be maintained that completely and accurately describes the TCB in terms of exceptions, error messages, and effects. A formal top-level specification (FTLS) of the TCB shall be maintained that accurately describes the TCB in terms of exceptions, error messages, and effects. The DTLS and FTLS shall include those components of the TCB that are implemented as hardware and/or firmware if their properties are visible at the TCB interface. The FTLS shall be shown to be an accurate description of the TCB interface. A convincing argument shall be given that the DTLS is consistent with the model and a combination of formal and informal techniques shall be used to show that the FTLS is consistent with the model. This verification evidence shall be consistent with that provided within the state-of-the-art of the particular computer security center-endorsed formal specification and verification system used. Manual or other mapping of the FTLS to the TCB source code shall be performed to provide evidence of correct implementation.

4.1.3.2.3 Configuration Management

During the entire life-cycle, i.e., during the design, development, and maintenance of the TCB, a configuration management system shall be in place for all security-relevant hardware, firmware, and software that maintains control of changes to the formal model, the descriptive and formal top-level specifications, other design data, implementation documentation, source code, the running version of the object code, and test fixtures and documentation. The configuration management system shall assure a consistent mapping among all documentation and code associated with the current version of the TCB. Tools shall be provided for generation of a new version of the TCB from source code. Also available shall be tools, maintained under strict configuration control, for comparing a newly generated version with the previous TCB version in order to ascertain that only the intended changes have been made in the code that will actually be used as the new version of the TCB. A combination of technical, physical, and procedural safeguards shall be used to protect from unauthorized modification or destruction the master copy or copies of all material used to generate the TCB.

4.1.3.2.4 Trusted Distribution

A trusted ADP system control and distribution facility shall be provided for maintaining the integrity of the mapping between the master data describing the current version of the TCB and the on-site master copy of the code for the current version. Procedures (e.g., site security acceptance testing) shall exist for assuring that the TCB software, firmware, and hardware updates distributed to a customer are exactly as specified by the master copies.

4.1.4 Documentation

4.1.4.1 Security Features User's Guide

A single summary, chapter, or manual in user documentation shall describe the protection mechanisms provided by the TCB, guidelines on their use, and how they interact with one another.

4.1.4.2 Trusted Facility Manual

A manual addressed to the ADP system administrator shall present cautions about functions and privileges that should be controlled when running a secure facility. The procedures for examining and maintaining the audit files as well as the detailed audit record structure for each type of audit event shall be given. The manual shall describe the operator and administrator functions related to security, to include changing the security characteristics of a user. It shall provide guidelines on the consistent and effective use of the protection features of the system, how they interact, how to securely generate a new TCB, and facility procedures, warnings, and privileges that need to be controlled in order to operate the facility in a secure manner. The TCB modules that contain the reference validation mechanism shall be identified. The procedures for secure generation of a new TCB from source after modification of any modules in the TCB shall be described. It shall include the procedures to ensure that the system is initially started in a secure manner. Procedures shall also be included to resume secure system operation after any lapse in system operation.

4.1.4.3 Test Documentation

The system developer shall provide to the evaluators a document that describes the test plan, test procedures that show how the security mechanisms were tested, and results of the security mechanisms' functional testing. It shall include results of testing the effectiveness of the methods used to reduce covert channel bandwidths. The results of the mapping between the formal top-level specification and the TCB source code shall be given.

4.1.4.4 Design Documentation

Documentation shall be available that provides a description of the manufacturer's philosophy of protection and an explanation of how this philosophy is translated into the TCB. The interfaces between the TCB modules shall be described. A formal description of the security policy model enforced by the TCB shall be available and proven that it is sufficient to enforce the security policy. The specific TCB protection mechanisms shall be identified and an explanation given to show that they satisfy the model. The descriptive top-level specification (DTLS) shall be shown to be an accurate description of the TCB interface. Documentation shall describe how the TCB implements the reference monitor concept and give an explanation why it is tamper resistant, cannot be bypassed, and is correctly implemented. The TCB implementation (i.e., in hardware, firmware, and software) shall be informally shown to be consistent with the formal top-level specification (FTLS). The elements of the FTLS shall be shown, using informal techniques, to correspond to the elements of the TCB. Documentation shall describe how the TCB is structured to facilitate testing and to enforce least privilege. This documentation shall also present the results of the covert channel analysis and the tradeoffs involved in restricting the channels. All auditable events that may be used in the exploitation of known covert storage channels shall be identified. The bandwidths of known covert storage channels, the use of which is not detectable by the auditing mechanisms, shall be provided. (See the Covert Channel Guideline section.) Hardware, firmware, and software mechanisms not dealt with in the FTLS but strictly internal to the TCB (e.g., mapping registers, direct memory access I/O) shall be clearly described.

4.2 BEYOND CLASS (A1)

Most of the security enhancements envisioned for systems that will provide features and assurance in addition to that already provided by class (A1) systems are beyond current technology. The discussion below is intended to guide future work and is derived from research and development activities already underway in both the public and private sectors. As more and better analysis techniques are developed, the requirements for these systems will become more explicit. In the future, use of formal verification will be extended to the source level and covert timing channels will be more fully addressed. At this level the design environment will become important and testing will be aided by analysis of the formal top-level specification. Consideration will be given to the correctness of the tools used in TCB development (e.g., compilers, assemblers, loaders) and to the correct functioning of the hardware/firmware on which the TCB will run. Areas to be addressed by systems beyond class (A1) include:

- * System Architecture

A demonstration (formal or otherwise) must be given showing that requirements of self-protection and completeness for reference monitors have been implemented in the TCB.

- * Security Testing

Although beyond the current state-of-the-art, it is envisioned that some test-case generation will be done automatically from the formal top-level specification or formal lower-level specifications.

- * Formal Specification and Verification

The TCB must be verified down to the source code level, using formal verification methods where feasible. Formal verification of the source code of the security-relevant portions of an operating system has proven to be a difficult task. Two important considerations are the choice of a high-level language whose semantics can be fully and formally expressed, and a careful mapping, through successive stages, of the abstract formal design to a formalization of the implementation in low-level specifications. Experience has shown that only when the lowest level specifications closely correspond to the actual code can code proofs be successfully accomplished.

- * Trusted Design Environment

The TCB must be designed in a trusted facility with only trusted (cleared) personnel.

PART II:
RATIONALE AND GUIDELINES

5.0 CONTROL OBJECTIVES FOR TRUSTED COMPUTER SYSTEMS

The criteria are divided within each class into groups of requirements. These groupings were developed to assure that three basic control objectives for computer security are satisfied and not overlooked. These control objectives deal with:

- * Security Policy
- * Accountability
- * Assurance

This section provides a discussion of these general control objectives and their implication in terms of designing trusted systems.

5.1 A NEED FOR CONSENSUS

A major goal of the DoD Computer Security Center is to encourage the Computer Industry to develop trusted computer systems and products, making them widely available in the commercial market place. Achievement of this goal will require recognition and articulation by both the public and private sectors of a need and demand for such products.

As described in the introduction to this document, efforts to define the problems and develop solutions associated with processing nationally sensitive information, as well as other sensitive data such as financial, medical, and personnel information used by the National Security Establishment, have been underway for a number of years. The criteria, as described in Part I, represent the culmination of these efforts and describe basic requirements for building trusted computer systems. To date, however, these systems have been viewed by many as only satisfying National Security needs. As long as this perception continues the consensus needed to motivate manufacture of trusted systems will be lacking.

The purpose of this section is to describe in detail the fundamental control objectives. These objectives lay the foundation for the requirements outlined in the criteria. The goal is to explain the foundations so that those outside the National Security Establishment can assess their universality and, by extension, the universal applicability of the criteria requirements to processing all types of sensitive applications whether they be for National Security or the private sector.

5.2 DEFINITION AND USEFULNESS

The term "control objective" refers to a statement of intent with respect to control over some aspect of an organization's resources, or processes, or both. In terms of a computer system, control objectives provide a framework for developing a strategy for fulfilling a set of security requirements for any given system. Developed in response to generic vulnerabilities, such as the need to manage and handle sensitive data in order to prevent compromise, or the need to provide accountability in order to detect fraud, control objectives have been identified as a useful method of expressing security goals.[3]

Examples of control objectives include the three basic design requirements for implementing the reference monitor concept discussed in Section 6. They are:

- * The reference validation mechanism must be tamperproof.
- * The reference validation mechanism must always be invoked.
- * The reference validation mechanism must be small enough to be subjected to analysis and tests, the completeness of which can be assured.[1]

5.3 CRITERIA CONTROL OBJECTIVES

The three basic control objectives of the criteria are concerned with security policy, accountability, and assurance. The remainder of this section provides a discussion of these basic requirements.

5.3.1 Security Policy

In the most general sense, computer security is concerned with controlling the way in which a computer can be used, i.e., controlling how information processed by it can be accessed and manipulated. However, at closer examination, computer security can refer to a number of areas. Symptomatic of this, FIPS PUB 39, Glossary For Computer Systems Security, does not have a unique definition for computer security.[16] Instead there are eleven separate definitions for security which include: ADP systems security, administrative security, data security, etc. A common thread running through these definitions is the word "protection." Further declarations of protection requirements can be found in DoD Directive 5200.28 which describes an acceptable level of protection for classified data to be one that will "assure that systems which process, store, or use classified data and produce classified information will, with reasonable dependability, prevent: a. Deliberate or inadvertent access to classified material by unauthorized persons, and b. Unauthorized manipulation of the computer and its associated peripheral devices." [8]

In summary, protection requirements must be defined in terms of the perceived threats, risks, and goals of an organization. This is often stated in terms of a security policy. It has been pointed out in the literature that it is external laws, rules, regulations, etc. that establish what access to information is to be permitted, independent of the use of a computer. In particular, a given system can only be said to be secure with respect to its enforcement of some specific policy.[30] Thus, the control objective for security policy is:

SECURITY POLICY CONTROL OBJECTIVE

A statement of intent with regard to control over access to and dissemination of information, to be known as the security policy must be precisely defined and implemented for each system that is used to process sensitive information. The security policy must accurately reflect the laws, regulations, and general policies from which it is derived.

5.3.1.1 Mandatory Security Policy

Where a security policy is developed that is to be applied to control of classified or other specifically designated sensitive information, the policy must include detailed rules on how to handle that information throughout its

life-cycle. These rules are a function of the various sensitivity designations that the information can assume and the various forms of access supported by the system. Mandatory security refers to the enforcement of a set of access control rules that constrains a subject's access to information on the basis of a comparison of that individual's clearance/authorization to the information, the classification/sensitivity designation of the information, and the form of access being mediated. Mandatory policies either require or can be satisfied by systems that can enforce a partial ordering of designations, namely, the designations must form what is mathematically known as a "lattice." [5]

A clear implication of the above is that the system must assure that the designations associated with sensitive data cannot be arbitrarily changed, since this could permit individuals who lack the appropriate authorization to access sensitive information. Also implied is the requirement that the system control the flow of information so that data cannot be stored with lower sensitivity designations unless its "downgrading" has been authorized. The control objective is:

MANDATORY SECURITY CONTROL OBJECTIVE

Security policies defined for systems that are used to process classified or other specifically categorized sensitive information must include provisions for the enforcement of mandatory access control rules. That is, they must include a set of rules for controlling access based directly on a comparison of the individual's clearance or authorization for the information and the classification or sensitivity designation of the information being sought, and indirectly on considerations of physical and other environmental factors of control. The mandatory access control rules must accurately reflect the laws, regulations, and general policies from which they are derived.

5.3.1.2 Discretionary Security Policy

Discretionary security is the principal type of access control available in computer systems today. The basis of this kind of security is that an individual user, or program operating on his behalf, is allowed to specify explicitly the types of access other users may have to information under his control. Discretionary security differs from mandatory security in that it implements an access control policy on the basis of an individual's need-to-know as opposed to mandatory controls which are driven by the classification or sensitivity designation of the information.

Discretionary controls are not a replacement for mandatory controls. In an environment in which information is classified (as in the DoD) discretionary security provides for a finer granularity of control within the overall constraints of the mandatory policy. Access to classified information requires effective implementation of both types of controls as precondition to granting that access. In general, no person may have access to classified information unless: (a) that person has been determined to be trustworthy, i.e., granted a personnel security clearance -- MANDATORY, and (b) access is necessary for the performance of official duties, i.e., determined to have a need-to-know -- DISCRETIONARY. In other words, discretionary controls give individuals discretion to decide on which of the permissible accesses will actually be allowed to which users, consistent with overriding mandatory policy restrictions. The control objective is:

DISCRETIONARY SECURITY CONTROL OBJECTIVE

Security policies defined for systems that are used to process classified or other sensitive information must include provisions for the enforcement of discretionary access control rules. That is, they must include a consistent set of rules for controlling and limiting access based on identified individuals who have been determined to have a need-to-know for the information.

5.3.1.3 Marking

To implement a set of mechanisms that will put into effect a mandatory security policy, it is necessary that the system mark information with appropriate classification or sensitivity labels and maintain these markings as the information moves through the system. Once information is unalterably and accurately marked, comparisons required by the mandatory access control rules can be accurately and consistently made. An additional benefit of having the system maintain the classification or sensitivity label internally is the ability to automatically generate properly "labeled" output. The labels, if accurately and integrally maintained by the system, remain accurate when output from the system. The control objective is:

MARKING CONTROL OBJECTIVE

Systems that are designed to enforce a mandatory security policy must store and preserve the integrity of classification or other sensitivity labels for all information. Labels exported from the system must be accurate representations of the corresponding internal sensitivity labels being exported.

5.3.2 Accountability

The second basic control objective addresses one of the fundamental principles of security, i.e., individual accountability. Individual accountability is the key to securing and controlling any system that processes information on behalf of individuals or groups of individuals. A number of requirements must be met in order to satisfy this objective.

The first requirement is for individual user identification. Second, there is a need for authentication of the identification. Identification is functionally dependent on authentication. Without authentication, user identification has no credibility. Without a credible identity, neither mandatory nor discretionary security policies can be properly invoked because there is no assurance that proper authorizations can be made.

The third requirement is for dependable audit capabilities. That is, a trusted computer system must provide authorized personnel with the ability to audit any action that can potentially cause access to, generation of, or effect the release of classified or sensitive information. The audit data will be selectively acquired based on the auditing needs of a particular installation and/or application. However, there must be sufficient granularity in the audit data to support tracing the auditable events to a specific individual who has taken the actions or on whose behalf the actions were taken. The control objective is:

ACCOUNTABILITY CONTROL OBJECTIVE

Systems that are used to process or handle classified or other sensitive information must assure individual accountability whenever either a mandatory or discretionary security policy is invoked. Furthermore, to assure accountability, the capability must exist for an authorized and competent agent to access and evaluate accountability information by a secure means, within a reasonable amount of time, and without undue difficulty.

5.3.3 Assurance

The third basic control objective is concerned with guaranteeing or providing confidence that the security policy has been implemented correctly and that the protection-relevant elements of the system do, indeed, accurately mediate and enforce the intent of that policy. By extension, assurance must include a guarantee that the trusted portion of the system works only as intended. To accomplish these objectives, two types of assurance are needed. They are life-cycle assurance and operational assurance.

Life-cycle assurance refers to steps taken by an organization to ensure that the system is designed, developed, and maintained using formalized and rigorous controls and standards.[17] Computer systems that process and store sensitive or classified information depend on the hardware and software to protect that information. It follows that the hardware and software themselves

must be protected against unauthorized changes that could cause protection mechanisms to malfunction or be bypassed completely. For this reason trusted computer systems must be carefully evaluated and tested during the design and development phases and reevaluated whenever changes are made that could affect the integrity of the protection mechanisms. Only in this way can confidence be provided that the hardware and software interpretation of the security policy is maintained accurately and without distortion.

While life-cycle assurance is concerned with procedures for managing system design, development, and maintenance; operational assurance focuses on features and system architecture used to ensure that the security policy is uncircumventably enforced during system operation. That is, the security policy must be integrated into the hardware and software protection features of the system. Examples of steps taken to provide this kind of confidence include: methods for testing the operational hardware and software for correct operation, isolation of protection-critical code, and the use of hardware and software to provide distinct domains. The control objective is:

ASSURANCE CONTROL OBJECTIVE

Systems that are used to process or handle classified or other sensitive information must be designed to guarantee correct and accurate interpretation of the security policy and must not distort the intent of that policy. Assurance must be provided that correct implementation and operation of the policy exists throughout the system's life-cycle.

6.0 RATIONALE BEHIND THE EVALUATION CLASSES

6.1 THE REFERENCE MONITOR CONCEPT

In October of 1972, the Computer Security Technology Planning Study, conducted by James P. Anderson & Co., produced a report for the Electronic Systems Division (ESD) of the United States Air Force.[1] In that report, the concept of "a reference monitor which enforces the authorized access relationships between subjects and objects of a system" was introduced. The reference monitor concept was found to be an essential element of any system that would provide multilevel secure computing facilities and controls.

The Anderson report went on to define the reference validation mechanism as "an implementation of the reference monitor concept . . . that validates each reference to data or programs by any user (program) against a list of authorized types of reference for that user." It then listed the three design requirements that must be met by a reference validation mechanism:

- a. The reference validation mechanism must be tamper proof.
- b. The reference validation mechanism must always be invoked.
- c. The reference validation mechanism must be small enough to be subject to analysis and tests, the completeness of which can be assured."[1]

Extensive peer review and continuing research and development activities have sustained the validity of the Anderson Committee's findings. Early examples of the reference validation mechanism were known as security kernels. The Anderson Report described the security kernel as "that combination of hardware and software which implements the reference monitor concept." [1] In this vein, it will be noted that the security kernel must support the three reference monitor requirements listed above.

6.2 A FORMAL SECURITY POLICY MODEL

Following the publication of the Anderson report, considerable research was initiated into formal models of security policy requirements and of the mechanisms that would implement and enforce those policy models as a security kernel. Prominent among these efforts was the ESD-sponsored development of the Bell and LaPadula model, an abstract formal treatment of DoD security policy.[2] Using mathematics and set theory, the model precisely defines the notion of secure state, fundamental modes of access, and the rules for granting subjects specific modes of access to objects. Finally, a theorem is proven to demonstrate that the rules are security-preserving operations, so that the application of any sequence of the rules to a system that is in a secure state will result in the system entering a new state that is also secure. This theorem is known as the Basic Security Theorem.

A subject can act on behalf of a user or another subject. The subject is created as a surrogate for the cleared user and is assigned a formal security level based on their classification. The state transitions and invariants of the formal policy model define the invariant relationships that must hold

between the clearance of the user, the formal security level of any process that can act on the user's behalf, and the formal security level of the devices and other objects to which any process can obtain specific modes of access. The Bell and LaPadula model, for example, defines a relationship between formal security levels of subjects and objects, now referenced as the "dominance relation." From this definition, accesses permitted between subjects and objects are explicitly defined for the fundamental modes of access, including read-only access, read/write access, and write-only access. The model defines the Simple Security Condition to control granting a subject read access to a specific object, and the *-Property (read "Star Property") to control granting a subject write access to a specific object. Both the Simple Security Condition and the *-Property include mandatory security provisions based on the dominance relation between formal security levels of subjects and objects the clearance of the subject and the classification of the object. The Discretionary Security Property is also defined, and requires that a specific subject be authorized for the particular mode of access required for the state transition. In its treatment of subjects (processes acting on behalf of a user), the model distinguishes between trusted subjects (i.e., not constrained within the model by the *-Property) and untrusted subjects (those that are constrained by the *-Property).

From the Bell and LaPadula model there evolved a model of the method of proof required to formally demonstrate that all arbitrary sequences of state transitions are security-preserving. It was also shown that the *-Property is sufficient to prevent the compromise of information by Trojan Horse attacks.

6.3 THE TRUSTED COMPUTING BASE

In order to encourage the widespread commercial availability of trusted computer systems, these evaluation criteria have been designed to address those systems in which a security kernel is specifically implemented as well as those in which a security kernel has not been implemented. The latter case includes those systems in which objective (c) is not fully supported because of the size or complexity of the reference validation mechanism. For convenience, these evaluation criteria use the term Trusted Computing Base to refer to the reference validation mechanism, be it a security kernel, front-end security filter, or the entire trusted computer system.

The heart of a trusted computer system is the Trusted Computing Base (TCB) which contains all of the elements of the system responsible for supporting the security policy and supporting the isolation of objects (code and data) on which the protection is based. The bounds of the TCB equate to the "security perimeter" referenced in some computer security literature. In the interest of understandable and maintainable protection, a TCB should be as simple as possible consistent with the functions it has to perform. Thus, the TCB includes hardware, firmware, and software critical to protection and must be designed and implemented such that system elements excluded from it need not be trusted to maintain protection. Identification of the interface and elements of the TCB along with their correct functionality therefore forms the basis for evaluation.

For general-purpose systems, the TCB will include key elements of the operating system and may include all of the operating system. For embedded systems, the security policy may deal with objects in a way that is meaningful at the application level rather than at the operating system level. Thus, the protection policy may be enforced in the application software rather than in the underlying operating system. The TCB will necessarily include all those portions of the operating system and application software essential to the support of the policy. Note that, as the amount of code in the TCB increases, it becomes harder to be confident that the TCB enforces the reference monitor requirements under all circumstances.

6.4 ASSURANCE

The third reference monitor design objective is currently interpreted as meaning that the TCB "must be of sufficiently simple organization and complexity to be subjected to analysis and tests, the completeness of which can be assured."

Clearly, as the perceived degree of risk increases (e.g., the range of sensitivity of the system's protected data, along with the range of clearances held by the system's user population) for a particular system's operational application and environment, so also must the assurances be increased to substantiate the degree of trust that will be placed in the system. The hierarchy of requirements that are presented for the evaluation classes in the trusted computer system evaluation criteria reflect the need for these assurances.

As discussed in Section 5.3, the evaluation criteria uniformly require a statement of the security policy that is enforced by each trusted computer system. In addition, it is required that a convincing argument be presented that explains why the TCB satisfies the first two design requirements for a reference monitor. It is not expected that this argument will be entirely formal. This argument is required for each candidate system in order to satisfy the assurance control objective.

The systems to which security enforcement mechanisms have been added, rather than built-in as fundamental design objectives, are not readily amenable to extensive analysis since they lack the requisite conceptual simplicity of a security kernel. This is because their TCB extends to cover much of the entire system. Hence, their degree of trustworthiness can best be ascertained only by obtaining test results. Since no test procedure for something as complex as a computer system can be truly exhaustive, there is always the possibility that a subsequent penetration attempt could succeed. It is for this reason that such systems must fall into the lower evaluation classes.

On the other hand, those systems that are designed and engineered to support the TCB concepts are more amenable to analysis and structured testing. Formal methods can be used to analyze the correctness of their reference validation mechanisms in enforcing the system's security policy. Other methods, including less-formal arguments, can be used in order to substantiate claims for the completeness of their access mediation and their degree of

tamper-resistance. More confidence can be placed in the results of this analysis and in the thoroughness of the structured testing than can be placed in the results for less methodically structured systems. For these reasons, it appears reasonable to conclude that these systems could be used in higher-risk environments. Successful implementations of such systems would be placed in the higher evaluation classes.

6.5 THE CLASSES

It is highly desirable that there be only a small number of overall evaluation classes. Three major divisions have been identified in the evaluation criteria with a fourth division reserved for those systems that have been evaluated and found to offer unacceptable security protection. Within each major evaluation division, it was found that "intermediate" classes of trusted system design and development could meaningfully be defined. These intermediate classes have been designated in the criteria because they identify systems that:

- * are viewed to offer significantly better protection and assurance than would systems that satisfy the basic requirements for their evaluation class; and
- * there is reason to believe that systems in the intermediate evaluation classes could eventually be evolved such that they would satisfy the requirements for the next higher evaluation class.

Except within division A it is not anticipated that additional "intermediate" evaluation classes satisfying the two characteristics described above will be identified.

Distinctions in terms of system architecture, security policy enforcement, and evidence of credibility between evaluation classes have been defined such that the "jump" between evaluation classes would require a considerable investment of effort on the part of implementors. Correspondingly, there are expected to be significant differentials of risk to which systems from the higher evaluation classes will be exposed.

7.0 THE RELATIONSHIP BETWEEN POLICY AND THE CRITERIA

Section 1 presents fundamental computer security requirements and Section 5 presents the control objectives for Trusted Computer Systems. They are general requirements, useful and necessary, for the development of all secure systems. However, when designing systems that will be used to process classified or other sensitive information, functional requirements for meeting the Control Objectives become more specific. There is a large body of policy laid down in the form of Regulations, Directives, Presidential Executive Orders, and OMB Circulars that form the basis of the procedures for the handling and processing of Federal information in general and classified information specifically. This section presents pertinent excerpts from these policy statements and discusses their relationship to the Control Objectives. These excerpts are examples to illustrate the relationship of the policies to criteria and may not be complete.

7.1 ESTABLISHED FEDERAL POLICIES

A significant number of computer security policies and associated requirements have been promulgated by Federal government elements. The interested reader is referred to reference [32] which analyzes the need for trusted systems in the civilian agencies of the Federal government, as well as in state and local governments and in the private sector. This reference also details a number of relevant Federal statutes, policies and requirements not treated further below.

Security guidance for Federal automated information systems is provided by the Office of Management and Budget. Two specifically applicable Circulars have been issued. OMB Circular No. A-71, Transmittal Memorandum No. 1, "Security of Federal Automated Information Systems,"[26] directs each executive agency to establish and maintain a computer security program. It makes the head of each executive branch, department and agency responsible "for assuring an adequate level of security for all agency data whether processed in-house or commercially. This includes responsibility for the establishment of physical, administrative and technical safeguards required to adequately protect personal, proprietary or other sensitive data not subject to national security regulations, as well as national security data."[26, para. 4 p. 2]

OMB Circular No. A-123, "Internal Control Systems,"[27] issued to help eliminate fraud, waste, and abuse in government programs requires: (a) agency heads to issue internal control directives and assign responsibility, (b) managers to review programs for vulnerability, and (c) managers to perform periodic reviews to evaluate strengths and update controls. Soon after promulgation of OMB Circular A-123, the relationship of its internal control requirements to building secure computer systems was recognized.[4] While not stipulating computer controls specifically, the definition of Internal Controls in A-123 makes it clear that computer systems are to be included:

"Internal Controls - The plan of organization and all of the methods and measures adopted within an agency to safeguard its resources, assure the accuracy and reliability of its information, assure adherence to applicable laws, regulations and policies, and promote operational economy and efficiency."[27, sec. 4.C]

The matter of classified national security information processed by ADP systems was one of the first areas given serious and extensive concern in computer security. The computer security policy documents promulgated as a result contain generally more specific and structured requirements than most, keyed in turn to an authoritative basis that itself provides a rather clearly articulated and structured information security policy. This basis, Executive Order 12356, "National Security Information," sets forth requirements for the classification, declassification and safeguarding of "national security information" per se.[14]

7.2 DOD POLICIES

Within the Department of Defense, these broad requirements are implemented and further specified primarily through two vehicles: 1) DoD Regulation 5200.1-R

[7], which applies to all components of the DoD as such, and 2) DoD 5220.22-M, "Industrial Security Manual for Safeguarding Classified Information" [11], which applies to contractors included within the Defense Industrial Security Program. Note that the latter transcends DoD as such, since it applies not only to any contractors handling classified information for any DoD component, but also to the contractors of eighteen other Federal organizations for whom the Secretary of Defense is authorized to act in rendering industrial security services.*

For ADP systems, these information security requirements are further amplified and specified in: 1) DoD Directive 5200.28 [8] and DoD Manual 5200.28-M [9], for DoD components; and 2) Section XIII of DoD 5220.22-M [11] for contractors. DoD Directive 5200.28, "Security Requirements for Automatic Data Processing (ADP) Systems," stipulates: "Classified material contained in an ADP system shall be safeguarded by the continuous employment of protective features in the system's hardware and software design and configuration" [8, sec. IV] Furthermore, it is required that ADP systems that "process, store, or use classified data and produce classified information will, with reasonable dependability, prevent:

- a. Deliberate or inadvertent access to classified material by unauthorized persons, and
- b. Unauthorized manipulation of the computer and its associated peripheral devices." [8, sec. I B.3]

Requirements equivalent to these appear within DoD 5200.28-M [9] and in DoD 5220.22-M [11].

DoD Directive 5200.28 provides the security requirements for ADP systems. For some types of information, such as Sensitive Compartmented Information (SCI), DoD Directive 5200.28 states that other minimum security requirements also apply. These minima are found in DCID 1/16 (new reference number 5) which is implemented in DIAM 50-4 (new reference number 6) for DoD and DoD contractor ADP systems.

From requirements imposed by these regulations, directives and circulars, the three components of the Security Policy Control Objective, i.e., Mandatory and Discretionary Security and Marking, as well as the Accountability and Assurance Control Objectives, can be functionally defined for DoD applications. The following discussion provides further specificity in Policy for these Control Objectives.

* i.e., NASA, Commerce Department, GSA, State Department, Small Business Administration, National Science Foundation, Treasury Department, Transportation Department, Interior Department, Agriculture Department, U.S. Information Agency, Labor Department, Environmental Protection Agency, Justice Department, U.S. Arms Control and Disarmament Agency, Federal Emergency Management Agency, Federal Reserve System, and U.S. General Accounting Office.

7.3.1 Marking

7.3 CRITERIA CONTROL OBJECTIVE FOR SECURITY POLICY

The control objective for marking is: "Systems that are designed to enforce a mandatory security policy must store and preserve the integrity of classification or other sensitivity labels for all information. Labels exported from the system must be accurate representations of the corresponding internal sensitivity labels being exported."

DoD 5220.22-M, "Industrial Security Manual for Safeguarding Classified Information," explains in paragraph 11 the reasons for marking information:

"a. General. Classification designation by physical marking, notation or other means serves to warn and to inform the holder what degree of protection against unauthorized disclosure is required for that information or material." (14)

Marking requirements are given in a number of policy statements.

Executive Order 12356 (Sections 1.5.a and 1.5.a.1) requires that classification markings "shall be shown on the face of all classified documents, or clearly associated with other forms of classified information in a manner appropriate to the medium involved." [14]

DoD Regulation 5200.1-R (Section 1-500) requires that: ". . . information or material that requires protection against unauthorized disclosure in the interest of national security shall be classified in one of three designations, namely: 'Top Secret,' 'Secret' or 'Confidential.'" [7] (By extension, for use in computer processing, the unofficial designation "Unclassified" is used to indicate information that does not fall under one of the other three designations of classified information.)

DoD Regulation 5200.1-R (Section 4-304b) requires that: "ADP systems and word processing systems employing such media shall provide for internal classification marking to assure that classified information contained therein that is reproduced or generated, will bear applicable classification and associated markings." (This regulation provides for the exemption of certain existing systems where "internal classification and applicable associated markings cannot be implemented without extensive system modifications." [7] However, it is clear that future DoD ADP systems must be able to provide applicable and accurate labels for classified and other sensitive information.)

DoD Manual 5200.28-M (Section IV, 4-305d) requires the following: "Security Labels - All classified material accessible by or within

the ADP system shall be identified as to its security classification and access or dissemination limitations, and all output of the ADP system shall be appropriately marked."[9]

7.3.2 Mandatory Security

The control objective for mandatory security is: "Security policies defined for systems that are used to process classified or other specifically categorized sensitive information must include provisions for the enforcement of mandatory access control rules. That is, they must include a set of rules for controlling access based directly on a comparison of the individual's clearance or authorization for the information and the classification or sensitivity designation of the information being sought, and indirectly on considerations of physical and other environmental factors of control. The mandatory access control rules must accurately reflect the laws, regulations, and general policies from which they are derived."

There are a number of policy statements that are related to mandatory security.

Executive Order 12356 (Section 4.1.a) states that "a person is eligible for access to classified information provided that a determination of trustworthiness has been made by agency heads or designated officials and provided that such access is essential to the accomplishment of lawful and authorized Government purposes."[14]

DoD Regulation 5200.1-R (Chapter I, Section 3) defines a Special Access Program as "any program imposing 'need-to-know' or access controls beyond those normally provided for access to Confidential, Secret, or Top Secret information. Such a program includes, but is not limited to, special clearance, adjudication, or investigative requirements, special designation of officials authorized to determine 'need-to-know', or special lists of persons determined to have a 'need-to-know.'"[7, para. 1-328] This passage distinguishes between a 'discretionary' determination of need-to-know and formal need-to-know which is implemented through Special Access Programs. DoD Regulation 5200.1-R, paragraph 7-100 describes general requirements for trustworthiness (clearance) and need-to-know, and states that the individual with possession, knowledge or control of classified information has final responsibility for determining if conditions for access have been met. This regulation further stipulates that "no one has a right to have access to classified information solely by virtue of rank or position." [7, para. 7-100])

DoD Manual 5200.28-M (Section II 2-100) states that, "Personnel who develop, test (debug), maintain, or use programs which are classified or which will be used to access or develop classified material shall have a personnel security clearance and an access authorization (need-to-know), as appropriate for the highest classified and most restrictive category of classified material

which they will access under system constraints."[9]

DoD Manual 5220.22-M (Paragraph 3.a) defines access as "the ability and opportunity to obtain knowledge of classified information. An individual, in fact, may have access to classified information by being in a place where such information is kept, if the security measures which are in force do not prevent him from gaining knowledge of the classified information."[11]

The above mentioned Executive Order, Manual, Directives and Regulations clearly imply that a trusted computer system must assure that the classification labels associated with sensitive data cannot be arbitrarily changed, since this could permit individuals who lack the appropriate clearance to access classified information. Also implied is the requirement that a trusted computer system must control the flow of information so that data from a higher classification cannot be placed in a storage object of lower classification unless its "downgrading" has been authorized.

7.3.3 Discretionary Security

The term discretionary security refers to a computer system's ability to control information on an individual basis. It stems from the fact that even though an individual has all the formal clearances for access to specific classified information, each individual's access to information must be based on a demonstrated need-to-know. Because of this, it must be made clear that this requirement is not discretionary in a "take it or leave it" sense. The directives and regulations are explicit in stating that the need-to-know test must be satisfied before access can be granted to the classified information. The control objective for discretionary security is: "Security policies defined for systems that are used to process classified or other sensitive information must include provisions for the enforcement of discretionary access control rules. That is, they must include a consistent set of rules for controlling and limiting access based on identified individuals who have been determined to have a need-to-know for the information."

DoD Regulation 5200.1-R (Paragraph 7-100) In addition to excerpts already provided that touch on need-to-know, this section of the regulation stresses the need-to-know principle when it states "no person may have access to classified information unless . . . access is necessary for the performance of official duties."[7]

Also, DoD Manual 5220.22-M (Section III 20.a) states that "an individual shall be permitted to have access to classified information only . . . when the contractor determines that access is necessary in the performance of tasks or services essential to the fulfillment of a contract or program, i.e., the individual has a need-to-know."[11]

7.4 CRITERIA CONTROL OBJECTIVE FOR ACCOUNTABILITY

The control objective for accountability is: "Systems that are used to process or handle classified or other sensitive information must assure individual accountability whenever either a mandatory or discretionary security policy is invoked. Furthermore, to assure accountability the capability must exist for an authorized and competent agent to access and evaluate accountability information by a secure means, within a reasonable amount of time, and without undue difficulty."

This control objective is supported by the following citations:

DoD Directive 5200.28 (VI.A.1) states: "Each user's identity shall be positively established, and his access to the system, and his activity in the system (including material accessed and actions taken) controlled and open to scrutiny." [8]

DoD Manual 5200.28-M (Section V 5-100) states: "An audit log or file (manual, machine, or a combination of both) shall be maintained as a history of the use of the ADP System to permit a regular security review of system activity. (e.g., The log should record security related transactions, including each access to a classified file and the nature of the access, e.g., logins, production of accountable classified outputs, and creation of new classified files. Each classified file successfully accessed [regardless of the number of individual references] during each 'job' or 'interactive session' should also be recorded in the audit log. Much of the material in this log may also be required to assure that the system preserves information entrusted to it.)" [9]

DoD Manual 5200.28-M (Section IV 4-305f) states: "Where needed to assure control of access and individual accountability, each user or specific group of users shall be identified to the ADP System by appropriate administrative or hardware/software measures. Such identification measures must be in sufficient detail to enable the ADP System to provide the user only that material which he is authorized." [9]

DoD Manual 5200.28-M (Section I 1-102b) states:

"Component's Designated Approving Authorities, or their designees for this purpose . . . will assure:

.

(4) Maintenance of documentation on operating systems (O/S) and all modifications thereto, and its retention for a sufficient period of time to enable tracing of security-related defects to their point of origin or inclusion in the system.

.

(6) Establishment of procedures to discover, recover,

handle, and dispose of classified material improperly disclosed through system malfunction or personnel action.

(7) Proper disposition and correction of security deficiencies in all approved ADP Systems, and the effective use and disposition of system housekeeping or audit records, records of security violations or security-related system malfunctions, and records of tests of the security features of an ADP System."[9]

DoD Manual 5220.22-M (Section XIII 111) states: "Audit Trails

a. The general security requirement for any ADP system audit trail is that it provide a documented history of the use of the system. An approved audit trail will permit review of classified system activity and will provide a detailed activity record to facilitate reconstruction of events to determine the magnitude of compromise (if any) should a security malfunction occur. To fulfill this basic requirement, audit trail systems, manual, automated or a combination of both must document significant events occurring in the following areas of concern: (i) preparation of input data and dissemination of output data (i.e., reportable interactivity between users and system support personnel), (ii) activity involved within an ADP environment (e.g., ADP support personnel modification of security and related controls), and (iii) internal machine activity.

b. The audit trail for an ADP system approved to process classified information must be based on the above three areas and may be stylized to the particular system. All systems approved for classified processing should contain most if not all of the audit trail records listed below. The contractor's SPP documentation must identify and describe those applicable:

1. Personnel access;
2. Unauthorized and surreptitious entry into the central computer facility or remote terminal areas;
3. Start/stop time of classified processing indicating pertinent systems security initiation and termination events (e.g., upgrading/downgrading actions pursuant to paragraph 107);
4. All functions initiated by ADP system console operators;
5. Disconnects of remote terminals and peripheral devices (paragraph 107c);
6. Log-on and log-off user activity;

7. Unauthorized attempts to access files or programs, as well as all open, close, create, and file destroy actions;

8. Program aborts and anomalies including identification information (i.e., user/program name, time and location of incident, etc.);

9. System hardware additions, deletions and maintenance actions;

10. Generations and modifications affecting the security features of the system software.

c. The ADP system security supervisor or designee shall review the audit trail logs at least weekly to assure that all pertinent activity is properly recorded and that appropriate action has been taken to correct any anomaly. The majority of ADP systems in use today can develop audit trail systems in accord with the above; however, special systems such as weapons, communications, communications security, and tactical data exchange and display systems, may not be able to comply with all aspects of the above and may require individualized consideration by the cognizant security office.

d. Audit trail records shall be retained for a period of one inspection cycle."[11]

7.5 CRITERIA CONTROL OBJECTIVE FOR ASSURANCE

The control objective for assurance is: "Systems that are used to process or handle classified or other sensitive information must be designed to guarantee correct and accurate interpretation of the security policy and must not distort the intent of that policy. Assurance must be provided that correct implementation and operation of the policy exists throughout the system's life-cycle."

A basis for this objective can be found in the following sections of DoD Directive 5200.28:

DoD Directive 5200.28 (IV.B.1) stipulates: "Generally, security of an ADP system is most effective and economical if the system is designed originally to provide it. Each Department of Defense Component undertaking design of an ADP system which is expected to process, store, use, or produce classified material shall: From the beginning of the design process, consider the security policies, concepts, and measures prescribed in this Directive."[8]

DoD Directive 5200.28 (IV.C.5.a) states: "Provision may be made to permit adjustment of ADP system area controls to the level of protection required for the classification category and type(s) of material actually

being handled by the system, provided change procedures are developed and implemented which will prevent both the unauthorized access to classified material handled by the system and the unauthorized manipulation of the system and its components. Particular attention shall be given to the continuous protection of automated system security measures, techniques and procedures when the personnel security clearance level of users having access to the system changes."[8]

DoD Directive 5200.28 (VI.A.2) states: "Environmental Control. The ADP System shall be externally protected to minimize the likelihood of unauthorized access to system entry points, access to classified information in the system, or damage to the system."[8]

DoD Manual 5200.28-M (Section I 1-102b) states:

"Component's Designated Approving Authorities, or their designees for this purpose . . . will assure:

.

(5) Supervision, monitoring, and testing, as appropriate, of changes in an approved ADP System which could affect the security features of the system, so that a secure system is maintained.

.

(7) Proper disposition and correction of security deficiencies in all approved ADP Systems, and the effective use and disposition of system housekeeping or audit records, records of security violations or security-related system malfunctions, and records of tests of the security features of an ADP System.

(8) Conduct of competent system ST&E, timely review of system ST&E reports, and correction of deficiencies needed to support conditional or final approval or disapproval of an ADP System for the processing of classified information.

(9) Establishment, where appropriate, of a central ST&E coordination point for the maintenance of records of selected techniques, procedures, standards, and tests used in the testing and evaluation of security features of ADP Systems which may be suitable for validation and use by other Department of Defense Components."[9]

DoD Manual 5220.22-M (Section XIII 103a) requires: "the initial approval, in writing, of the cognizant security office prior to processing any classified information in an ADP system. This section requires reapproval by the cognizant security office for major system modifications made subsequent to initial approval. Reapprovals will be required because of (i) major changes in personnel access requirements, (ii) relocation or structural modification of the central computer facility, (iii) additions, deletions or changes to main frame, storage or

input/output devices, (iv) system software changes impacting security protection features, (v) any change in clearance, declassification, audit trail or hardware/software maintenance procedures, and (vi) other system changes as determined by the cognizant security office."[11]

A major component of assurance, life-cycle assurance, as described in DoD Directive 7920.1, is concerned with testing ADP systems both in the development phase as well as during operation (17). DoD Directive 5215.1 (Section F.2.C.(2)) requires "evaluations of selected industry and government-developed trusted computer systems against these criteria."[10]

8.0 A GUIDELINE ON COVERT CHANNELS

A covert channel is any communication channel that can be exploited by a process to transfer information in a manner that violates the system's security policy. There are two types of covert channels: storage channels and timing channels. Covert storage channels include all vehicles that would allow the direct or indirect writing of a storage location by one process and the direct or indirect reading of it by another. Covert timing channels include all vehicles that would allow one process to signal information to another process by modulating its own use of system resources in such a way that the change in response time observed by the second process would provide information.

From a security perspective, covert channels with low bandwidths represent a lower threat than those with high bandwidths. However, for many types of covert channels, techniques used to reduce the bandwidth below a certain rate (which depends on the specific channel mechanism and the system architecture) also have the effect of degrading the performance provided to legitimate system users. Hence, a trade-off between system performance and covert channel bandwidth must be made. Because of the threat of compromise that would be present in any multilevel computer system containing classified or sensitive information, such systems should not contain covert channels with high bandwidths. This guideline is intended to provide system developers with an idea of just how high a "high" covert channel bandwidth is.

A covert channel bandwidth that exceeds a rate of one hundred (100) bits per second is considered "high" because 100 bits per second is the approximate rate at which many computer terminals are run. It does not seem appropriate to call a computer system "secure" if information can be compromised at a rate equal to the normal output rate of some commonly used device.

In any multilevel computer system there are a number of relatively low-bandwidth covert channels whose existence is deeply ingrained in the system design. Faced with the large potential cost of reducing the bandwidths of such covert channels, it is felt that those with maximum bandwidths of less than one (1) bit per second are acceptable in most application environments. Though maintaining acceptable performance in some systems may make it impractical to eliminate all covert channels with bandwidths of 1 or more bits per second, it is possible to audit their use without adversely affecting system performance. This audit capability provides the system administration with a means of detecting -- and procedurally correcting -- significant compromise. Therefore, a Trusted Computing Base should provide, wherever possible, the capability to audit the use of covert channel mechanisms with bandwidths that may exceed a rate of one (1) bit in ten (10) seconds.

The covert channel problem has been addressed by a number of authors. The interested reader is referred to references [5], [6], [19], [21], [22], [23], and [29].

9.0 A GUIDELINE ON CONFIGURING MANDATORY ACCESS CONTROL FEATURES

The Mandatory Access Control requirement includes a capability to support an unspecified number of hierarchical classifications and an unspecified number of non-hierarchical categories at each hierarchical level. To encourage consistency and portability in the design and development of the National Security Establishment trusted computer systems, it is desirable for all such systems to be able to support a minimum number of levels and categories. The following suggestions are provided for this purpose:

- * The number of hierarchical classifications should be greater than or equal to sixteen (16).
- * The number of non-hierarchical categories should be greater than or equal to sixty-four (64).

10.0 A GUIDELINE ON SECURITY TESTING

These guidelines are provided to give an indication of the extent and sophistication of testing undertaken by the DoD Computer Security Center during the Formal Product Evaluation process. Organizations wishing to use "Department of Defense Trusted Computer System Evaluation Criteria" for performing their own evaluations may find this section useful for planning purposes.

As in Part I, highlighting is used to indicate changes in the guidelines from the next lower division.

10.1 TESTING FOR DIVISION C

10.1.1 Personnel

The security testing team shall consist of at least two individuals with bachelor degrees in Computer Science or the equivalent. Team members shall be able to follow test plans prepared by the system developer and suggest additions, shall be familiar with the "flaw hypothesis" or equivalent security testing methodology, and shall have assembly level programming experience. Before testing begins, the team members shall have functional knowledge of, and shall have completed the system developer's internals course for, the system being evaluated.

10.1.2 Testing

The team shall have "hands-on" involvement in an independent run of the tests used by the system developer. The team shall independently design and implement at least five system-specific tests in an attempt to circumvent the security mechanisms of the system. The elapsed time devoted to testing shall be at least one month and need not exceed three months. There shall be no fewer than twenty hands-on hours spent carrying out system developer-defined tests and test team-defined tests.

10.2 TESTING FOR DIVISION B

10.2.1 Personnel

The security testing team shall consist of at least two individuals with bachelor degrees in Computer Science or the equivalent and at least one individual with a master's degree in Computer Science or equivalent. Team members shall be able to follow test plans prepared by the system developer and suggest additions, shall be conversant with the "flaw hypothesis" or equivalent security testing methodology, shall be fluent in the TCB implementation language(s), and shall have assembly level programming experience. Before testing begins, the team members shall have functional knowledge of, and shall have completed the system developer's internals course for, the system being evaluated. At least one team member shall have previously completed a security test on another system.

10.2.2 Testing

The team shall have "hands-on" involvement in an independent run of the test package used by the system developer to test security-relevant hardware and software. The team shall independently design and implement at least fifteen system-specific tests in an attempt to circumvent the security mechanisms of the system. The elapsed time devoted to testing

shall be at least two months and need not exceed four months. There shall be no fewer than thirty hands-on hours per team member spent carrying out system developer-defined tests and test team-defined tests.

10.3 TESTING FOR DIVISION A

10.3.1 Personnel

The security testing team shall consist of at least one individual with a bachelor's degree in Computer Science or the equivalent and at least two individuals with masters' degrees in Computer Science or equivalent. Team members shall be able to follow test plans prepared by the system developer and suggest additions, shall be conversant with the "flaw hypothesis" or equivalent security testing methodology, shall be fluent in the TCB implementation language(s), and shall have assembly level programming experience. Before testing begins, the team members shall have functional knowledge of, and shall have completed the system developer's internals course for, the system being evaluated. At least one team member shall be familiar enough with the system hardware to understand the maintenance diagnostic programs and supporting hardware documentation. At least two team members shall have previously completed a security test on another system. At least one team member shall have demonstrated system level programming competence on the system under test to a level of complexity equivalent to adding a device driver to the system.

10.3.2 Testing

The team shall have "hands-on" involvement in an independent run of the test package used by the system developer to test security-relevant hardware and software. The team shall independently design and implement at least twenty-five system-specific tests in an attempt to circumvent the security mechanisms of the system. The elapsed time devoted to testing shall be at least three months and need not exceed six months. There shall be no fewer than fifty hands-on hours per team member spent carrying out system developer-defined tests and test team-defined tests.

APPENDIX A

COMMERCIAL PRODUCE EVALUATION PROCESS

"Department of Defense Trusted Computer System Evaluation Criteria" forms the basis upon which the Computer Security Center will carry out the commercial computer security evaluation process. This process is focused on commercially produced and supported general-purpose operating system products that meet the needs of government departments and agencies. The formal evaluation is aimed at "off-the-shelf" commercially supported products and is completely divorced from any consideration of overall system performance, potential applications, or particular processing environments. The evaluation provides a key input to a computer system security approval/accreditation. However, it does not constitute a complete computer system security evaluation. A complete study (e.g., as in reference [18]) must consider additional factors dealing with the system in its unique environment, such as its proposed security mode of operation, specific users, applications, data sensitivity, physical and personnel security, administrative and procedural security, TEMPEST, and communications security.

The product evaluation process carried out by the Computer Security Center has three distinct elements:

- * Preliminary Product Evaluation - An informal dialogue between a vendor and the Center in which technical information is exchanged to create a common understanding of the vendor's product, the criteria, and the rating that may be expected to result from a formal product evaluation.
- * Formal Product Evaluation - A formal evaluation, by the Center, of a product that is available to the DoD, and that results in that product and its assigned rating being placed on the Evaluated Products List.
- * Evaluated Products List - A list of products that have been subjected to formal product evaluation and their assigned ratings.

Preliminary Product Evaluation

Since it is generally very difficult to add effective security measures late in a product's life cycle, the Center is interested in working with system vendors in the early stages of product design. A preliminary product evaluation allows the Center to consult with computer vendors on computer security issues found in products that have not yet been formally announced.

A preliminary evaluation is typically initiated by computer system vendors who are planning new computer products that feature security or major security-related upgrades to existing products. After an initial meeting between the vendor and the Center, appropriate non-disclosure agreements are executed that require the Center to maintain the confidentiality of any proprietary information disclosed to it. Technical exchange meetings follow in which the vendor provides details about the proposed product (particularly its internal designs and goals) and the Center provides expert feedback to the

vendor on potential computer security strengths and weaknesses of the vendor's design choices, as well as relevant interpretation of the criteria. The preliminary evaluation is typically terminated when the product is completed and ready for field release by the vendor. Upon termination, the Center prepares a wrap-up report for the vendor and for internal distribution within the Center. Those reports containing proprietary information are not available to the public.

During preliminary evaluation, the vendor is under no obligation to actually complete or market the potential product. The Center is, likewise, not committed to conduct a formal product evaluation. A preliminary evaluation may be terminated by either the Center or the vendor when one notifies the other, in writing, that it is no longer advantageous to continue the evaluation.

Formal Product Evaluation

The formal product evaluation provides a key input to certification of a computer system for use in National Security Establishment applications and is the sole basis for a product being placed on the Evaluated Products List.

A formal product evaluation begins with a request by a vendor for the Center to evaluate a product for which the product itself and accompanying documentation needed to meet the requirements defined by this publication are complete. Non-disclosure agreements are executed and a formal product evaluation team is formed by the Center. An initial meeting is then held with the vendor to work out the schedule for the formal evaluation. Since testing of the implemented product forms an important part of the evaluation process, access by the evaluation team to a working version of the system is negotiated with the vendor. Additional support required from the vendor includes complete design documentation, source code, and access to vendor personnel who can answer detailed questions about specific portions of the product. The evaluation team tests the product against each requirement, making any necessary interpretations of the criteria with respect to the product being evaluated.

The evaluation team writes a final report on their findings about the system. The report is publicly available (containing no proprietary or sensitive information) and contains the overall class rating assigned to the system and the details of the evaluation team's findings when comparing the product against the evaluation criteria. Detailed information concerning vulnerabilities found by the evaluation team is furnished to the system developers and designers as each is found so that the vendor has a chance to eliminate as many of them as possible prior to the completion of the Formal Product Evaluation. Vulnerability analyses and other proprietary or sensitive information are controlled within the Center through the Vulnerability Reporting Program and are distributed only within the U.S. Government on a strict need-to-know and non-disclosure basis, and to the vendor.

APPENDIX B

SUMMARY OF EVALUATION CRITERIA DIVISIONS

The divisions of systems recognized under the trusted computer system evaluation criteria are as follows. Each division represents a major improvement in the overall confidence one can place in the system to protect classified and other sensitive information.

Division (D): Minimal Protection

This division contains only one class. It is reserved for those systems that have been evaluated but that fail to meet the requirements for a higher evaluation class.

Division (C): Discretionary Protection

Classes in this division provide for discretionary (need-to-know) protection and, through the inclusion of audit capabilities, for accountability of subjects and the actions they initiate.

Division (B): Mandatory Protection

The notion of a TCB that preserves the integrity of sensitivity labels and uses them to enforce a set of mandatory access control rules is a major requirement in this division. Systems in this division must carry the sensitivity labels with major data structures in the system. The system developer also provides the security policy model on which the TCB is based and furnishes a specification of the TCB. Evidence must be provided to demonstrate that the reference monitor concept has been implemented.

Division (A): Verified Protection

This division is characterized by the use of formal security verification methods to assure that the mandatory and discretionary security controls employed in the system can effectively protect classified or other sensitive information stored or processed by the system. Extensive documentation is required to demonstrate that the TCB meets the security requirements in all aspects of design, development and implementation.

APPENDIX C

SUMMARY OF EVALUATION CRITERIA CLASSES

The classes of systems recognized under the trusted computer system evaluation criteria are as follows. They are presented in the order of increasing desirability from a computer security point of view.

Class (D): Minimal Protection

This class is reserved for those systems that have been evaluated but that fail to meet the requirements for a higher evaluation class.

Class (C1): Discretionary Security Protection

The Trusted Computing Base (TCB) of a class (C1) system nominally satisfies the discretionary security requirements by providing separation of users and data. It incorporates some form of credible controls capable of enforcing access limitations on an individual basis, i.e., ostensibly suitable for allowing users to be able to protect project or private information and to keep other users from accidentally reading or destroying their data. The class (C1) environment is expected to be one of cooperating users processing data at the same level(s) of sensitivity.

Class (C2): Controlled Access Protection

Systems in this class enforce a more finely grained discretionary access control than (C1) systems, making users individually accountable for their actions through login procedures, auditing of security-relevant events, and resource isolation.

Class (B1): Labeled Security Protection

Class (B1) systems require all the features required for class (C2). In addition, an informal statement of the security policy model, data labeling, and mandatory access control over named subjects and objects must be present. The capability must exist for accurately labeling exported information. Any flaws identified by testing must be removed.

Class (B2): Structured Protection

In class (B2) systems, the TCB is based on a clearly defined and documented formal security policy model that requires the discretionary and mandatory access control enforcement found in class (B1) systems be extended to all subjects and objects in the ADP system. In addition, covert channels are addressed. The TCB must be carefully structured into protection-critical and non- protection-critical elements. The TCB interface is well-defined and the TCB design and implementation enable it to be subjected to more thorough testing and more complete review. Authentication mechanisms are strengthened, trusted facility management is provided in the form of support for system administrator and operator functions, and stringent configuration management controls are imposed. The system is relatively resistant to penetration.

Class (B3): Security Domains

The class (B3) TCB must satisfy the reference monitor requirements that it mediate all accesses of subjects to objects, be tamperproof, and be small enough to be subjected to analysis and tests. To this end, the TCB is structured to exclude code not essential to security policy enforcement, with significant system engineering during TCB design and implementation directed toward minimizing its complexity. A security administrator is supported, audit mechanisms are expanded to signal security- relevant events, and system recovery procedures are required. The system is highly resistant to penetration.

Class (A1): Verified Design

Systems in class (A1) are functionally equivalent to those in class (B3) in that no additional architectural features or policy requirements are added. The distinguishing feature of systems in this class is the analysis derived from formal design specification and verification techniques and the resulting high degree of assurance that the TCB is correctly implemented. This assurance is developmental in nature, starting with a formal model of the security policy and a formal top-level specification (FTLS) of the design. In keeping with the extensive design and development analysis of the TCB required of systems in class (A1), more stringent configuration management is required and procedures are established for securely distributing the system to sites. A system security administrator is supported.

APPENDIX D

REQUIREMENT DIRECTORY

This appendix lists requirements defined in "Department of Defense Trusted Computer System Evaluation Criteria" alphabetically rather than by class. It is provided to assist in following the evolution of a requirement through the classes. For each requirement, three types of criteria may be present. Each will be preceded by the word: NEW, CHANGE, or ADD to indicate the following:

NEW: Any criteria appearing in a lower class are superseded by the criteria that follow.

CHANGE: The criteria that follow have appeared in a lower class but are changed for this class. Highlighting is used to indicate the specific changes to previously stated criteria.

ADD: The criteria that follow have not been required for any lower class, and are added in this class to the previously stated criteria for this requirement.

Abbreviations are used as follows:

NR: (No Requirement) This requirement is not included in this class.

NAR: (No Additional Requirements) This requirement does not change from the previous class.

The reader is referred to Part I of this document when placing new criteria for a requirement into the complete context for that class.

Figure 1 provides a pictorial summary of the evolution of requirements through the classes.

Audit

C1: NR.

C2: NEW: The TCB shall be able to create, maintain, and protect from modification or unauthorized access or destruction an audit trail of accesses to the objects it protects. The audit data shall be protected by the TCB so that read access to it is limited to those who are authorized for audit data. The TCB shall be able to record the following types of events: use of identification and authentication mechanisms, introduction of objects into a user's address space (e.g., file open, program initiation), deletion of objects, and actions taken by computer operators and system administrators and/or system security officers and other security relevant events. For each recorded event, the audit record shall

identify: date and time of the event, user, type of event, and success or failure of the event. For identification/authentication events the origin of request (e.g., terminal ID) shall be included in the audit record. For events that introduce an object into a user's address space and for object deletion events the audit record shall include the name of the object. The ADP system administrator shall be able to selectively audit the actions of any one or more users based on individual identity.

B1: CHANGE: For events that introduce an object into a user's address space and for object deletion events the audit record shall include the name of the object and the object's security level. The ADP system administrator shall be able to selectively audit the actions of any one or more users based on individual identity and/or object security level.

ADD: The TCB shall also be able to audit any override of human-readable output markings.

B2: ADD: The TCB shall be able to audit the identified events that may be used in the exploitation of covert storage channels.

B3: ADD: The TCB shall contain a mechanism that is able to monitor the occurrence or accumulation of security auditable events that may indicate an imminent violation of security policy. This mechanism shall be able to immediately notify the security administrator when thresholds are exceeded, and, if the occurrence or accumulation of these security relevant events continues, the system shall take the least disruptive action to terminate the event.

A1: NAR.

Configuration Management

C1: NR.

C2: NR.

B1: NR.

B2: NEW: During development and maintenance of the TCB, a configuration management system shall be in place that maintains control of changes to the descriptive top-level specification, other design data, implementation documentation, source code, the running version of the object code, and test fixtures and documentation. The configuration management system shall assure a consistent mapping among all documentation and code associated with the current version of the TCB. Tools shall be provided for generation of a new version of the TCB from source code. Also available shall be tools for comparing a newly generated version with the previous TCB version in order to ascertain that only the intended changes have been made in the code that will actually be used as the new version of the TCB.

B3: NAR.

A1: CHANGE: During the entire life-cycle, i.e., during the design, development, and maintenance of the TCB, a configuration management system shall be in place for all security-relevant hardware, firmware, and software that maintains control of changes to the formal model, the descriptive and formal top-level specifications, other design data, implementation documentation, source code, the running version of the object code, and test fixtures and documentation. Also available shall be tools, maintained under strict configuration control, for comparing a newly generated version with the previous TCB version in order to ascertain that only the intended changes have been made in the code that will actually be used as the new version of the TCB.

ADD: A combination of technical, physical, and procedural safeguards shall be used to protect from unauthorized modification or destruction the master copy or copies of all material used to generate the TCB.

Covert Channel Analysis

C1: NR.

C2: NR.

B1: NR.

B2: NEW: The system developer shall conduct a thorough search for covert storage channels and make a determination (either by actual measurement or by engineering estimation) of the maximum bandwidth of each identified channel. (See the Covert Channels Guideline section.)

B3: CHANGE: The system developer shall conduct a thorough search for covert channels and make a determination (either by actual measurement or by engineering estimation) of the maximum bandwidth of each identified channel.

A1: ADD: Formal methods shall be used in the analysis.

Design Documentation

C1: NEW: Documentation shall be available that provides a description of the manufacturer's philosophy of protection and an explanation of how this philosophy is translated into the TCB. If the TCB is composed of distinct modules, the interfaces between these modules shall be described.

C2: NAR.

B1: ADD: An informal or formal description of the security policy model enforced by the TCB shall be available and an explanation provided to show that it is sufficient to enforce the security policy. The specific TCB protection mechanisms shall be identified and an

explanation given to show that they satisfy the model.

B2: CHANGE: The interfaces between the TCB modules shall be described. A formal description of the security policy model enforced by the TCB shall be available and proven that it is sufficient to enforce the security policy.

ADD: The descriptive top-level specification (DTLS) shall be shown to be an accurate description of the TCB interface. Documentation shall describe how the TCB implements the reference monitor concept and give an explanation why it is tamper resistant, cannot be bypassed, and is correctly implemented. Documentation shall describe how the TCB is structured to facilitate testing and to enforce least privilege. This documentation shall also present the results of the covert channel analysis and the tradeoffs involved in restricting the channels. All auditable events that may be used in the exploitation of known covert storage channels shall be identified. The bandwidths of known covert storage channels, the use of which is not detectable by the auditing mechanisms, shall be provided. (See the Covert Channel Guideline section.)

B3: ADD: The TCB implementation (i.e., in hardware, firmware, and software) shall be informally shown to be consistent with the DTLS. The elements of the DTLS shall be shown, using informal techniques, to correspond to the elements of the TCB.

A1: CHANGE: The TCB implementation (i.e., in hardware, firmware, and software) shall be informally shown to be consistent with the formal top-level specification (FTLS). The elements of the FTLS shall be shown, using informal techniques, to correspond to the elements of the TCB.

ADD: Hardware, firmware, and software mechanisms not dealt with in the FTLS but strictly internal to the TCB (e.g., mapping registers, direct memory access I/O) shall be clearly described.

Design Specification and Verification

C1: NR.

C2: NR.

B1: NEW: An informal or formal model of the security policy supported by the TCB shall be maintained over the life cycle of the ADP system that is shown to be consistent with its axioms.

B2: CHANGE: A formal model of the security policy supported by the TCB shall be maintained over the life cycle of the ADP system that is proven consistent with its axioms.

ADD: A descriptive top-level specification (DTLS) of the TCB shall be maintained that completely and accurately describes the TCB in terms of exceptions, error messages, and effects. It shall be shown to be an accurate description of the TCB interface.

B3: ADD: A convincing argument shall be given that the DTLS is consistent with the model.

A1: CHANGE: The FTLS shall be shown to be an accurate description of the TCB interface. A convincing argument shall be given that the DTLS is consistent with the model and a combination of formal and informal techniques shall be used to show that the FTLS is consistent with the model.

ADD: A formal top-level specification (FTLS) of the TCB shall be maintained that accurately describes the TCB in terms of exceptions, error messages, and effects. The DTLS and FTLS shall include those components of the TCB that are implemented as hardware and/or firmware if their properties are visible at the TCB interface. This verification evidence shall be consistent with that provided within the state-of-the-art of the particular Computer Security Center-endorsed formal specification and verification system used. Manual or other mapping of the FTLS to the TCB source code shall be performed to provide evidence of correct implementation.

Device Labels

C1: NR.

C2: NR.

B1: NR.

B2: NEW: The TCB shall support the assignment of minimum and maximum security levels to all attached physical devices. These security levels shall be used by the TCB to enforce constraints imposed by the physical environments in which the devices are located.

B3: NAR.

A1: NAR.

Discretionary Access Control

C1: NEW: The TCB shall define and control access between named users and named objects (e.g., files and programs) in the ADP system. The enforcement mechanism (e.g., self/group/public controls, access control lists) shall allow users to specify and control sharing of those objects by named individuals or defined groups or both.

C2: CHANGE: The enforcement mechanism (e.g., self/group/public controls, access control lists) shall allow users to specify and control sharing of those objects by named individuals, or defined groups of individuals, or by both, and shall provide controls to limit propagation of access rights.

ADD: The discretionary access control mechanism shall, either by explicit user action or by default, provide that objects are protected from

unauthorized access. These access controls shall be capable of including or excluding access to the granularity of a single user. Access permission to an object by users not already possessing access permission shall only be assigned by authorized users.

B1: NAR.

B2: NAR.

B3: CHANGE: The enforcement mechanism (e.g., access control lists) shall allow users to specify and control sharing of those objects, and shall provide controls to limit propagation of access rights. These access controls shall be capable of specifying, for each named object, a list of named individuals and a list of groups of named individuals with their respective modes of access to that object.

ADD: Furthermore, for each such named object, it shall be possible to specify a list of named individuals and a list of groups of named individuals for which no access to the object is to be given.

A1: NAR.

Exportation of Labeled Information

C1: NR.

C2: NR.

B1: NEW: The TCB shall designate each communication channel and I/O device as either single-level or multilevel. Any change in this designation shall be done manually and shall be auditable by the TCB. The TCB shall maintain and be able to audit any change in the security level or levels associated with a communication channel or I/O device.

B2: NAR.

B3: NAR.

A1: NAR.

Exportation to Multilevel Devices

C1: NR.

C2: NR.

B1: NEW: When the TCB exports an object to a multilevel I/O device, the sensitivity label associated with that object shall also be exported and shall reside on the same physical medium as the exported information and shall be in the same form (i.e., machine-readable or human-readable form). When the TCB exports or imports an object over a multilevel communication channel, the protocol used on that channel shall provide for the unambiguous pairing between the sensitivity

labels and the associated information that is sent or received.

B2: NAR.

B3: NAR.

A1: NAR.

Exportation to Single-Level Devices

C1: NR.

C2: NR.

B1: NEW: Single-level I/O devices and single-level communication channels are not required to maintain the sensitivity labels of the information they process. However, the TCB shall include a mechanism by which the TCB and an authorized user reliably communicate to designate the single security level of information imported or exported via single-level communication channels or I/O devices.

B2: NAR.

B3: NAR.

A1: NAR.

Identification and Authentication

C1: NEW: The TCB shall require users to identify themselves to it before beginning to perform any other actions that the TCB is expected to mediate. Furthermore, the TCB shall use a protected mechanism (e.g., passwords) to authenticate the user's identity. The TCB shall protect authentication data so that it cannot be accessed by any unauthorized user.

C2: ADD: The TCB shall be able to enforce individual accountability by providing the capability to uniquely identify each individual ADP system user. The TCB shall also provide the capability of associating this identity with all auditable actions taken by that individual.

B1: CHANGE: Furthermore, the TCB shall maintain authentication data that includes information for verifying the identity of individual users (e.g., passwords) as well as information for determining the clearance and authorizations of individual users. This data shall be used by the TCB to authenticate the user's identity and to ensure that the security level and authorizations of subjects external to the TCB that may be created to act on behalf of the individual user are dominated by the clearance and authorization of that user.

B2: NAR.

B3: NAR.

A1: NAR.

Label Integrity

C1: NR.

C2: NR.

B1: NEW: Sensitivity labels shall accurately represent security levels of the specific subjects or objects with which they are associated. When exported by the TCB, sensitivity labels shall accurately and unambiguously represent the internal labels and shall be associated with the information being exported.

B2: NAR.

B3: NAR.

A1: NAR.

Labeling Human-Readable Output

C1: NR.

C2: NR.

B1: NEW: The ADP system administrator shall be able to specify the printable label names associated with exported sensitivity labels. The TCB shall mark the beginning and end of all human-readable, paged, hardcopy output (e.g., line printer output) with human-readable sensitivity labels that properly* represent the sensitivity of the output. The TCB shall, by default, mark the top and bottom of each page of human-readable, paged, hardcopy output (e.g., line printer output) with human-readable sensitivity labels that properly* represent the overall sensitivity of the output or that properly* represent the sensitivity of the information on the page. The TCB shall, by default and in an appropriate manner, mark other forms of human-readable output (e.g., maps, graphics) with human-readable sensitivity labels that properly* represent the sensitivity of the output. Any override of these marking defaults shall be auditable by the TCB.

B2: NAR.

B3: NAR.

* The hierarchical classification component in human-readable sensitivity labels shall be equal to the greatest hierarchical classification of any of the information in the output that the labels refer to; the non-hierarchical category component shall include all of the non-hierarchical categories of the information in the output the labels refer to, but no other non-hierarchical

categories.

A1: NAR.

Labels

C1: NR.

C2: NR.

B1: NEW: Sensitivity labels associated with each subject and storage object under its control (e.g., process, file, segment, device) shall be maintained by the TCB. These labels shall be used as the basis for mandatory access control decisions. In order to import non-labeled data, the TCB shall request and receive from an authorized user the security level of the data, and all such actions shall be auditable by the TCB.

B2: CHANGE: Sensitivity labels associated with each ADP system resource (e.g., subject, storage object, ROM) that is directly or indirectly accessible by subjects external to the TCB shall be maintained by the TCB.

B3: NAR.

A1: NAR.

Mandatory Access Control

C1: NR.

C2: NR.

B1: NEW: The TCB shall enforce a mandatory access control policy over all subjects and storage objects under its control (e.g., processes, files, segments, devices). These subjects and objects shall be assigned sensitivity labels that are a combination of hierarchical classification levels and non-hierarchical categories, and the labels shall be used as the basis for mandatory access control decisions. The TCB shall be able to support two or more such security levels. (See the Mandatory Access Control guidelines.) The following requirements shall hold for all accesses between subjects and objects controlled by the TCB: A subject can read an object only if the hierarchical classification in the subject's security level is greater than or equal to the hierarchical classification in the object's security level and the non-hierarchical categories in the subject's security level include all the non-hierarchical categories in the object's security level. A subject can write an object only if the hierarchical classification in the subject's security level is less than or equal to the hierarchical classification in the object's security level and all the non-hierarchical categories in the subject's security level are included in the non-hierarchical categories in the object's security level. Identification and

authentication data shall be used by the TCB to authenticate the user's identity and to ensure that the security level and authorization of subjects external to the TCB that may be created to act on behalf of the individual user are dominated by the clearance and authorization of that user.

B2: CHANGE: The TCB shall enforce a mandatory access control policy over all resources (i.e., subjects, storage objects, and I/O devices) that are directly or indirectly accessible by subjects external to the TCB. The following requirements shall hold for all accesses between all subjects external to the TCB and all objects directly or indirectly accessible by these subjects:

B3: NAR.

A1: NAR.

Object Reuse

C1: NR.

C2: NEW: All authorizations to the information contained within a storage object shall be revoked prior to initial assignment, allocation or reallocation to a subject from the TCB's pool of unused storage objects. No information, including encrypted representations of information, produced by a prior subject's actions is to be available to any subject that obtains access to an object that has been released back to the system.

B1: NAR.

B2: NAR.

B3: NAR.

A1: NAR.

Security Features User's Guide

C1: NEW: A single summary, chapter, or manual in user documentation shall describe the protection mechanisms provided by the TCB, guidelines on their use, and how they interact with one another.

C2: NAR.

B1: NAR.

B2: NAR.

B3: NAR.

A1: NAR.

Security Testing

- C1: NEW: The security mechanisms of the ADP system shall be tested and found to work as claimed in the system documentation. Testing shall be done to assure that there are no obvious ways for an unauthorized user to bypass or otherwise defeat the security protection mechanisms of the TCB. (See the Security Testing guidelines.)
- C2: ADD: Testing shall also include a search for obvious flaws that would allow violation of resource isolation, or that would permit unauthorized access to the audit or authentication data.
- B1: NEW: The security mechanisms of the ADP system shall be tested and found to work as claimed in the system documentation. A team of individuals who thoroughly understand the specific implementation of the TCB shall subject its design documentation, source code, and object code to thorough analysis and testing. Their objectives shall be: to uncover all design and implementation flaws that would permit a subject external to the TCB to read, change, or delete data normally denied under the mandatory or discretionary security policy enforced by the TCB; as well as to assure that no subject (without authorization to do so) is able to cause the TCB to enter a state such that it is unable to respond to communications initiated by other users. All discovered flaws shall be removed or neutralized and the TCB retested to demonstrate that they have been eliminated and that new flaws have not been introduced. (See the Security Testing Guidelines.)
- B2: CHANGE: All discovered flaws shall be corrected and the TCB retested to demonstrate that they have been eliminated and that new flaws have not been introduced.
- ADD: The TCB shall be found relatively resistant to penetration. Testing shall demonstrate that the TCB implementation is consistent with the descriptive top-level specification.
- B3: CHANGE: The TCB shall be found resistant to penetration.
- ADD: No design flaws and no more than a few correctable implementation flaws may be found during testing and there shall be reasonable confidence that few remain.
- A1: CHANGE: Testing shall demonstrate that the TCB implementation is consistent with the formal top-level specification.
- ADD: Manual or other mapping of the FTLS to the source code may form a basis for penetration testing.

Subject Sensitivity Labels

C1: NR.

C2: NR.

B1: NR.

B2: NEW: The TCB shall immediately notify a terminal user of each change in the security level associated with that user during an interactive session. A terminal user shall be able to query the TCB as desired for a display of the subject's complete sensitivity label.

B3: NAR.

A1: NAR.

System Architecture

C1: NEW: The TCB shall maintain a domain for its own execution that protects it from external interference or tampering (e.g., by modification of its code or data structures). Resources controlled by the TCB may be a defined subset of the subjects and objects in the ADP system.

C2: ADD: The TCB shall isolate the resources to be protected so that they are subject to the access control and auditing requirements.

B1: ADD: The TCB shall maintain process isolation through the provision of distinct address spaces under its control.

B2: NEW: The TCB shall maintain a domain for its own execution that protects it from external interference or tampering (e.g., by modification of its code or data structures). The TCB shall maintain process isolation through the provision of distinct address spaces under its control. The TCB shall be internally structured into well-defined largely independent modules. It shall make effective use of available hardware to separate those elements that are protection-critical from those that are not. The TCB modules shall be designed such that the principle of least privilege is enforced. Features in hardware, such as segmentation, shall be used to support logically distinct storage objects with separate attributes (namely: readable, writeable). The user interface to the TCB shall be completely defined and all elements of the TCB identified.

B3: ADD: The TCB shall be designed and structured to use a complete, conceptually simple protection mechanism with precisely defined semantics. This mechanism shall play a central role in enforcing the internal structuring of the TCB and the system. The TCB shall incorporate significant use of layering, abstraction and data hiding. Significant system engineering shall be directed toward minimizing the complexity of the TCB and excluding from the TCB modules that are not protection-critical.

A1: NAR.

System Integrity

C1: NEW: Hardware and/or software features shall be provided that can be used to periodically validate the correct operation of the on-site hardware and firmware elements of the TCB.

C2: NAR.

B1: NAR.

B2: NAR.

B3: NAR.

A1: NAR.

Test Documentation

C1: NEW: The system developer shall provide to the evaluators a document that describes the test plan, test procedures that show how the security mechanisms were tested and results of the security mechanisms' functional testing.

C2: NAR.

B1: NAR.

B2: ADD: It shall include results of testing the effectiveness of the methods used to reduce covert channel bandwidths.

B3: NAR.

A1: ADD: The results of the mapping between the formal top-level specification and the TCB source code shall be given.

Trusted Distribution

C1: NR.

C2: NR.

B1: NR.

B2: NR.

B3: NR.

A1: NEW: A trusted ADP system control and distribution facility shall be provided for maintaining the integrity of the mapping between the master data describing the current version of the TCB and the on-site master copy of the code for the current version. Procedures (e.g., site security acceptance testing) shall exist for assuring that the TCB software, firmware, and hardware updates distributed to a customer are exactly as specified by the master copies.

Trusted Facility Management

C1: NR.

C2: NR.

B1: NR.

B2: NEW: The TCB shall support separate operator and administrator functions.

B3: ADD: The functions performed in the role of a security administrator shall be identified. The ADP system administrative personnel shall only be able to perform security administrator functions after taking a distinct auditable action to assume the security administrator role on the ADP system. Non-security functions that can be performed in the security administration role shall be limited strictly to those essential to performing the security role effectively.

A1: NAR.

Trusted Facility Manual

C1: NEW: A manual addressed to the ADP system administrator shall present cautions about functions and privileges that should be controlled when running a secure facility.

C2: ADD: The procedures for examining and maintaining the audit files as well as the detailed audit record structure for each type of audit event shall be given.

B1: ADD: The manual shall describe the operator and administrator functions related to security, to include changing the characteristics of a user. It shall provide guidelines on the consistent and effective use of the protection features of the system, how they interact, how to securely generate a new TCB, and facility procedures, warnings, and privileges that need to be controlled in order to operate the facility in a secure manner.

B2: ADD: The TCB modules that contain the reference validation mechanism shall be identified. The procedures for secure generation of a new TCB from source after modification of any modules in the TCB shall be described.

B3: ADD: It shall include the procedures to ensure that the system is initially started in a secure manner. Procedures shall also be included to resume secure system operation after any lapse in system operation.

A1: NAR.

Trusted Path

C1: NR.

C2: NR.

B1: NR.

B2: NEW: The TCB shall support a trusted communication path between itself and user for initial login and authentication. Communications via this path shall be initiated exclusively by a user.

B3: CHANGE: The TCB shall support a trusted communication path between itself and users for use when a positive TCB-to-user connection is required (e.g., login, change subject security level). Communications via this trusted path shall be activated exclusively by a user or the TCB and shall be logically isolated and unmistakably distinguishable from other paths.

A1: NAR.

Trusted Recovery

C1: NR.

C2: NR.

B1: NR.

B2: NR.

B3: NEW: Procedures and/or mechanisms shall be provided to assure that, after an ADP system failure or other discontinuity, recovery without a protection compromise is obtained.

A1: NAR.

(this page is reserved for Figure 1)

GLOSSARY

Access - A specific type of interaction between a subject and an object that results in the flow of information from one to the other.

Approval/Accreditation - The official authorization that is granted to an ADP system to process sensitive information in its operational environment, based upon comprehensive security evaluation of the system's hardware, firmware, and software security design, configuration, and implementation and of the other system procedural, administrative, physical, TEMPEST, personnel, and communications security controls.

Audit Trail - A set of records that collectively provide documentary evidence of processing used to aid in tracing from original transactions forward to related records and reports, and/or backwards from records and reports to their component source transactions.

Authenticate - To establish the validity of a claimed identity.

Automatic Data Processing (ADP) System - An assembly of computer hardware, firmware, and software configured for the purpose of classifying, sorting, calculating, computing, summarizing, transmitting and receiving, storing, and retrieving data with a minimum of human intervention.

Bandwidth - A characteristic of a communication channel that is the amount of information that can be passed through it in a given amount of time, usually expressed in bits per second.

Bell-LaPadula Model - A formal state transition model of computer security policy that describes a set of access control rules. In this formal model, the entities in a computer system are divided into abstract sets of subjects and objects. The notion of a secure state is defined and it is proven that each state transition preserves security by moving from secure state to secure state; thus, inductively proving that the system is secure. A system state is defined to be "secure" if the only permitted access modes of subjects to objects are in accordance with a specific security policy. In order to determine whether or not a specific access mode is allowed, the clearance of a subject is compared to the classification of the object and a determination is made as to whether the subject is authorized for the specific access mode. The clearance/classification scheme is expressed in terms of a lattice. See also: Lattice, Simple Security Property, *-Property.

Certification - The technical evaluation of a system's security

features, made as part of and in support of the approval/accreditation process, that establishes the extent to which a particular computer system's design and implementation meet a set of specified security requirements.

Channel - An information transfer path within a system. May also refer to the mechanism by which the path is effected.

Covert Channel - A communication channel that allows a process to transfer information in a manner that violates the system's security policy. See also: Covert Storage Channel, Covert Timing Channel.

Covert Storage Channel - A covert channel that involves the direct or indirect writing of a storage location by one process and the direct or indirect reading of the storage location by another process. Covert storage channels typically involve a finite resource (e.g., sectors on a disk) that is shared by two subjects at different security levels.

Covert Timing Channel - A covert channel in which one process signals information to another by modulating its own use of system resources (e.g., CPU time) in such a way that this manipulation affects the real response time observed by the second process.

Data - Information with a specific physical representation.

Data Integrity - The state that exists when computerized data is the same as that in the source documents and has not been exposed to accidental or malicious alteration or destruction.

Descriptive Top-Level Specification (DTLS) - A top-level specification that is written in a natural language (e.g., English), an informal program design notation, or a combination of the two.

Discretionary Access Control - A means of restricting access to objects based on the identity of subjects and/or groups to which they belong. The controls are discretionary in the sense that a subject with a certain access permission is capable of passing that permission (perhaps indirectly) on to any other subject (unless restrained by mandatory access control).

Domain - The set of objects that a subject has the ability to access.

Dominate - Security level S1 is said to dominate security level S2 if the hierarchical classification of S1 is greater than or equal to that of S2 and the non-hierarchical categories

of S1 include all those of S2 as a subset.

Exploitable Channel - Any channel that is useable or detectable by subjects external to the Trusted Computing Base.

Flaw Hypothesis Methodology - A system analysis and penetration technique where specifications and documentation for the system are analyzed and then flaws in the system are hypothesized. The list of hypothesized flaws is then prioritized on the basis of the estimated probability that a flaw actually exists and, assuming a flaw does exist, on the ease of exploiting it and on the extent of control or compromise it would provide. The prioritized list is used to direct the actual testing of the system.

Flaw - An error of commission, omission, or oversight in a system that allows protection mechanisms to be bypassed.

Formal Proof - A complete and convincing mathematical argument, presenting the full logical justification for each proof step, for the truth of a theorem or set of theorems. The formal verification process uses formal proofs to show the truth of certain properties of formal specification and for showing that computer programs satisfy their specifications.

Formal Security Policy Model - A mathematically precise statement of a security policy. To be adequately precise, such a model must represent the initial state of a system, the way in which the system progresses from one state to another, and a definition of a "secure" state of the system. To be acceptable as a basis for a TCB, the model must be supported by a formal proof that if the initial state of the system satisfies the definition of a "secure" state and if all assumptions required by the model hold, then all future states of the system will be secure. Some formal modeling techniques include: state transition models, temporal logic models, denotational semantics models, algebraic specification models. An example is the model described by Bell and LaPadula in reference [2]. See also: Bell-LaPadula Model, Security Policy Model.

Formal Top-Level Specification (FTLS) - A Top-Level Specification that is written in a formal mathematical language to allow theorems showing the correspondence of the system specification to its formal requirements to be hypothesized and formally proven.

Formal Verification - The process of using formal proofs to demonstrate the consistency (design verification) between a formal specification of a system and a formal security policy model or (implementation verification) between the formal specification and its program implementation.

Front-End Security Filter - A process that is invoked to process

data according to a specified security policy prior to releasing the data outside the processing environment or upon receiving data from an external source.

Functional Testing - The portion of security testing in which the advertised features of a system are tested for correct operation.

General-Purpose System - A computer system that is designed to aid in solving a wide variety of problems.

Granularity - The relative fineness or coarseness by which a mechanism can be adjusted. The phrase "the granularity of a single user" means the access control mechanism can be adjusted to include or exclude any single user.

Lattice - A partially ordered set for which every pair of elements has a greatest lower bound and a least upper bound.

Least Privilege - This principle requires that each subject in a system be granted the most restrictive set of privileges (or lowest clearance) needed for the performance of authorized tasks. The application of this principle limits the damage that can result from accident, error, or unauthorized use.

Mandatory Access Control - A means of restricting access to objects based on the sensitivity (as represented by a label) of the information contained in the objects and the formal authorization (i.e., clearance) of subjects to access information of such sensitivity.

Multilevel Device - A device that is used in a manner that permits it to simultaneously process data of two or more security levels without risk of compromise. To accomplish this, sensitivity labels are normally stored on the same physical medium and in the same form (i.e., machine-readable or human-readable) as the data being processed.

Multilevel Secure - A class of system containing information with different sensitivities that simultaneously permits access by users with different security clearances and needs-to-know, but prevents users from obtaining access to information for which they lack authorization.

Object - A passive entity that contains or receives information. Access to an object potentially implies access to the information it contains. Examples of objects are: records, blocks, pages, segments, files, directories, directory trees, and programs, as well as bits, bytes, words, fields, processors, video displays, keyboards, clocks, printers, network nodes, etc.

Object Reuse - The reassignment to some subject of a medium (e.g., page frame, disk sector, magnetic tape) that

contained one or more objects. To be securely reassigned, such media must contain no residual data from the previously contained object(s).

Output - Information that has been exported by a TCB.

Password - A private character string that is used to authenticate an identity.

Penetration Testing - The portion of security testing in which the penetrators attempt to circumvent the security features of a system. The penetrators may be assumed to use all system design and implementation documentation, which may include listings of system source code, manuals, and circuit diagrams. The penetrators work under no constraints other than those that would be applied to ordinary users.

Process - A program in execution. It is completely characterized by a single current execution point (represented by the machine state) and address space.

Protection-Critical Portions of the TCB - Those portions of the TCB whose normal function is to deal with the control of access between subjects and objects.

Protection Philosophy - An informal description of the overall design of a system that delineates each of the protection mechanisms employed. A combination (appropriate to the evaluation class) of formal and informal techniques is used to show that the mechanisms are adequate to enforce the security policy.

Read - A fundamental operation that results only in the flow of information from an object to a subject.

Read Access - Permission to read information.

Read-Only Memory (ROM) - A storage area in which the contents can be read but not altered during normal computer processing.

Reference Monitor Concept - An access control concept that refers to an abstract machine that mediates all accesses to objects by subjects.

Resource - Anything used or consumed while performing a function. The categories of resources are: time, information, objects (information containers), or processors (the ability to use information). Specific examples are: CPU time; terminal connect time; amount of directly-addressable memory; disk space; number of I/O requests per minute, etc.

Security Kernel - The hardware, firmware, and software elements of a Trusted Computing Base that implement the reference monitor concept. It must mediate all accesses, be protected

from modification, and be verifiable as correct.

Security Level - The combination of a hierarchical classification and a set of non-hierarchical categories that represents the sensitivity of information.

Security Policy - The set of laws, rules, and practices that regulate how an organization manages, protects, and distributes sensitive information.

Security Policy Model - An informal presentation of a formal security policy model.

Security Relevant Event - Any event that attempts to change the security state of the system, (e.g., change discretionary access controls, change the security level of the subject, change user password, etc.). Also, any event that attempts to violate the security policy of the system, (e.g., too many attempts to login, attempts to violate the mandatory access control limits of a device, attempts to downgrade a file, etc.).

Security Testing - A process used to determine that the security features of a system are implemented as designed and that they are adequate for a proposed application environment. This process includes hands-on functional testing, penetration testing, and verification. See also: Functional Testing, Penetration Testing, Verification.

Sensitive Information - Information that, as determined by a competent authority, must be protected because its unauthorized disclosure, alteration, loss, or destruction will at least cause perceivable damage to someone or something.

Sensitivity Label - A piece of information that represents the security level of an object and that describes the sensitivity (e.g., classification) of the data in the object. Sensitivity labels are used by the TCB as the basis for mandatory access control decisions.

Simple Security Condition - A Bell-LaPadula security model rule allowing a subject read access to an object only if the security level of the subject dominates the security level of the object.

Single-Level Device - A device that is used to process data of a single security level at any one time. Since the device need not be trusted to separate data of different security levels, sensitivity labels do not have to be stored with the data being processed.

***-Property (Star Property)** - A Bell-LaPadula security model rule allowing a subject write access to an object only if the

security level of the subject is dominated by the security level of the object. Also known as the Confinement Property.

Storage Object - An object that supports both read and write accesses.

Subject - An active entity, generally in the form of a person, process, or device that causes information to flow among objects or changes the system state. Technically, a process/domain pair.

Subject Security Level - A subject's security level is equal to the security level of the objects to which it has both read and write access. A subject's security level must always be dominated by the clearance of the user the subject is associated with.

TEMPEST - The study and control of spurious electronic signals emitted from ADP equipment.

Top-Level Specification (TLS) - A non-procedural description of system behavior at the most abstract level. Typically a functional specification that omits all implementation details.

Trap Door - A hidden software or hardware mechanism that permits system protection mechanisms to be circumvented. It is activated in some non-apparent manner (e.g., special "random" key sequence at a terminal).

Trojan Horse - A computer program with an apparently or actually useful function that contains additional (hidden) functions that surreptitiously exploit the legitimate authorizations of the invoking process to the detriment of security. For example, making a "blind copy" of a sensitive file for the creator of the Trojan Horse.

Trusted Computer System - A system that employs sufficient hardware and software integrity measures to allow its use for processing simultaneously a range of sensitive or classified information.

Trusted Computing Base (TCB) - The totality of protection mechanisms within a computer system -- including hardware, firmware, and software -- the combination of which is responsible for enforcing a security policy. A TCB consists of one or more components that together enforce a unified security policy over a product or system. The ability of a trusted computing base to correctly enforce a security policy depends solely on the mechanisms within the TCB and on the correct input by system administrative personnel of parameters (e.g., a user's clearance) related to the security policy.

Trusted Path - A mechanism by which a person at a terminal can communicate directly with the Trusted Computing Base. This mechanism can only be activated by the person or the Trusted Computing Base and cannot be imitated by untrusted software.

Trusted Software - The software portion of a Trusted Computing Base.

User - Any person who interacts directly with a computer system.

Verification - The process of comparing two levels of system specification for proper correspondence (e.g., security policy model with top-level specification, TLS with source code, or source code with object code). This process may or may not be automated.

Write - A fundamental operation that results only in the flow of information from a subject to an object.

Write Access - Permission to write an object.

REFERENCES

1. Anderson, J. P. Computer Security Technology Planning Study, ESD-TR-73-51, vol. I, ESD/AFSC, Hanscom AFB, Bedford, Mass., October 1972 (NTIS AD-758 206).
2. Bell, D. E. and LaPadula, L. J. Secure Computer Systems: Unified Exposition and Multics Interpretation, MTR-2997 Rev. 1, MITRE Corp., Bedford, Mass., March 1976.
3. Brand, S. L. "An Approach to Identification and Audit of Vulnerabilities and Control in Application Systems," in Audit and Evaluation of Computer Security II: System Vulnerabilities and Controls, Z. Ruthberg, ed., NBS Special Publication #500-57, MD78733, April 1980.
4. Brand, S. L. "Data Processing and A-123," in Proceedings of the Computer Performance Evaluation User's Group 18th Meeting, C. B. Wilson, ed., NBS Special Publication #500-95, October 1982.
5. DCID 1/16, Security of Foreign Intelligence in Automated Data Processing Systems and Networks (U), 4 January 1983.
6. DIAM 50-4, Security of Compartmented Computer Operations (U), 24 June 1980.
7. Denning, D. E. "A Lattice Model of Secure Information Flow," in Communications of the ACM, vol. 19, no. 5 (May 1976), pp. 236-243.
8. Denning, D. E. Secure Information Flow in Computer Systems, Ph.D. dissertation, Purdue Univ., West Lafayette, Ind., May 1975.
9. DoD Directive 5000.29, Management of Computer Resources in Major Defense Systems, 26 April 1976.
10. DoD 5200.1-R, Information Security Program Regulation, August 1982.
11. DoD Directive 5200.28, Security Requirements for Automatic Data Processing (ADP) Systems, revised April 1978.
12. DoD 5200.28-M, ADP Security Manual -- Techniques and Procedures for Implementing, Deactivating, Testing, and Evaluating Secure Resource-Sharing ADP Systems, revised June 1979.
13. DoD Directive 5215.1, Computer Security Evaluation Center, 25 October 1982.

14. DoD 5220.22-M, Industrial Security Manual for Safeguarding Classified Information, March 1984.
15. DoD 5220.22-R, Industrial Security Regulation, February 1984.
16. DoD Directive 5400.11, Department of Defense Privacy Program, 9 June 1982.
17. DoD Directive 7920.1, Life Cycle Management of Automated Information Systems (AIS), 17 October 1978
18. Executive Order 12356, National Security Information, 6 April 1982.
19. Faurer, L. D. "Keeping the Secrets Secret," in Government Data Systems, November - December 1981, pp. 14-17.
20. Federal Information Processing Standards Publication (FIPS PUB) 39, Glossary for Computer Systems Security, 15 February 1976.
21. Federal Information Processing Standards Publication (FIPS PUB) 73, Guidelines for Security of Computer Applications, 30 June 1980.
22. Federal Information Processing Standards Publication (FIPS PUB) 102, Guideline for Computer Security Certification and Accreditation.
23. Lampson, B. W. "A Note on the Confinement Problem," in Communications of the ACM, vol. 16, no. 10 (October 1973), pp. 613-615.
24. Lee, T. M. P., et al. "Processors, Operating Systems and Nearby Peripherals: A Consensus Report," in Audit and Evaluation of Computer Security II: System Vulnerabilities and Controls, Z. Ruthberg, ed., NBS Special Publication #500-57, MD78733, April 1980.
25. Lipner, S. B. A Comment on the Confinement Problem, MITRE Corp., Bedford, Mass.
26. Millen, J. K. "An Example of a Formal Flow Violation," in Proceedings of the IEEE Computer Society 2nd International Computer Software and Applications Conference, November 1978, pp. 204-208.
27. Millen, J. K. "Security Kernel Validation in Practice," in Communications of the ACM, vol. 19, no. 5 (May 1976), pp. 243-250.
28. Nibaldi, G. H. Proposed Technical Evaluation Criteria for Trusted Computer Systems, MITRE Corp., Bedford, Mass., M79-225, AD-A108-832, 25 October 1979.

29. Nibaldi, G. H. Specification of A Trusted Computing Base, (TCB), MITRE Corp., Bedford, Mass., M79-228, AD-A108-831, 30 November 1979.
30. OMB Circular A-71, Transmittal Memorandum No. 1, Security of Federal Automated Information Systems, 27 July 1978.
31. OMB Circular A-123, Internal Control Systems, 5 November 1981.
32. Ruthberg, Z. and McKenzie, R., eds. Audit and Evaluation of Computer Security, in NBS Special Publication #500-19, October 1977.
33. Schaefer, M., Linde, R. R., et al. "Program Confinement in KVM/370," in Proceedings of the ACM National Conference, October 1977, Seattle.
34. Schell, R. R. "Security Kernels: A Methodical Design of System Security," in Technical Papers, USE Inc. Spring Conference, 5-9 March 1979, pp. 245-250.
35. Trotter, E. T. and Tasker, P. S. Industry Trusted Computer Systems Evaluation Process, MITRE Corp., Bedford, Mass., MTR-3931, 1 May 1980.
36. Turn, R. Trusted Computer Systems: Needs and Incentives for Use in government and Private Sector, (AD # A103399), Rand Corporation (R-28811-DR&E), June 1981.
37. Walker, S. T. "The Advent of Trusted Computer Operating Systems," in National Computer Conference Proceedings, May 1980, pp. 655-665.
38. Ware, W. H., ed., Security Controls for Computer Systems: Report of Defense Science Board Task Force on Computer Security, AD # A076617/0, Rand Corporation, Santa Monica, Calif., February 1970, reissued October 1979.