

FILE COPY

XRI Call No.

83342

Copy No. 1 of 2 cys.

ESD-TR-75-89

MTR-3024

A GENERAL SECURITY MARKING POLICY
FOR CLASSIFIED COMPUTER INPUT/OUTPUT MATERIAL

SEPTEMBER 1975

Prepared for

DEPUTY FOR COMMAND AND MANAGEMENT SYSTEMS
ELECTRONIC SYSTEMS DIVISION
AIR FORCE SYSTEMS COMMAND
UNITED STATES AIR FORCE
Hanscom Air Force Base, Bedford, Massachusetts



Approved for public release;
distribution unlimited.

Project No. 522B
Prepared by
THE MITRE CORPORATION
Bedford, Massachusetts
Contract No. F19628-75-C-0001

BEST AVAILABLE COPY


ADA016467


When U.S. Government drawings, specifications, or other data are used for any purpose other than a definitely related government procurement operation, the government thereby incurs no responsibility nor any obligation whatsoever; and the fact that the government may have formulated, furnished, or in any way supplied the said drawings, specifications, or other data is not to be regarded by implication or otherwise, as in any manner licensing the holder or any other person or corporation, or conveying any rights or permission to manufacture, use, or sell any patented invention that may in any way be related thereto.

Do not return this copy. Retain or destroy.


REVIEW AND APPROVAL

This technical report has been reviewed and is approved for publication.


PAUL A. KARGER, 1Lt, USAF
Techniques Engineering Division


THOMAS W. BAILEY, Major, USAF
Chief, Computer Security Branch

FOR THE COMMANDER


FRANK J. EMMA, Colonel, USAF
Director, Information Systems
Technology Applications Office
Deputy for Command and Management Systems

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER ESD-TR-75-89	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) A GENERAL SECURITY MARKING POLICY FOR CLASSIFIED INPUT/OUTPUT MATERIAL		5. TYPE OF REPORT & PERIOD COVERED
		6. PERFORMING ORG. REPORT NUMBER MTR-3024
7. AUTHOR(s) J. Mogilensky		8. CONTRACT OR GRANT NUMBER(s) F19628-75-C-0001
9. PERFORMING ORGANIZATION NAME AND ADDRESS The MITRE Corporation Box 208 Bedford, MA 01730		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS Project No. 522B
11. CONTROLLING OFFICE NAME AND ADDRESS Deputy for Command and Management Systems Electronic Systems Division, AFSC Hanscom Air Force Base, Bedford, MA 01731		12. REPORT DATE September 1975
		13. NUMBER OF PAGES 61
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report) UNCLASSIFIED
		15a. DECLASSIFICATION DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) COMPUTER OPERATIONS COMPUTER SECURITY INPUT/OUTPUT PROCESSING SECURITY MARKING POLICY		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) A security marking policy has two objectives: (1) to fulfill the basic marking requirements of the security regulations for all types of computer I/O material on a consistent, generalized basis, and (2) to ensure that the security attributes of classified data are accurately maintained on every I/O transfer. The reasons behind the need for a general policy are considered, the guidance provided by existing regulations is reviewed, and		

20. ABSTRACT

the marking policies of three existing systems are outlined. The basic concept of multilevel I/O is defined, a set of groups of I/O techniques is established, and two distinct marking policies, one for unilevel I/O and one for multilevel, are presented. Finally, an overview of the considerations involved in the implementation of trustworthy computer labeling is provided.

TABLE OF CONTENTS

	<u>Page</u>
LIST OF ILLUSTRATIONS	3
LIST OF TABLES	3
SECTION I INTRODUCTION	5
SECTION II THE NEED FOR A GENERAL POLICY	8
SECTION III GUIDANCE PROVIDED BY EXISTING REGULATIONS	10
INTRODUCTION	10
INFORMATION SECURITY PROGRAM REGULATION, DoD 5200.1-R and AFR 205-1	10
MANAGEMENT OF DATA PROCESSING EQUIPMENT, AFM 171-9	12
AIR FORCE MESSAGE MANAGEMENT PROGRAM, AFM 10-2	15
TECHNIQUES AND PROCEDURES FOR IMPLEMENT- ING, DEACTIVATING, TESTING, AND EVALUATING - SECURE RESOURCE-SHARING ADP SYSTEMS, DoD 5200.28-M	15
ABSTRACTION OF GENERAL PRINCIPLES	17
SECTION IV EXISTING POLICIES FOR SPECIFIC SYSTEMS	19
INTRODUCTION	19
MITRE CORPORATION COMPUTER FACILITY	19
WWMCCS GCOS-III SYSTEM	20
UNITED STATES AIR FORCE DATA SERVICES CENTER	22
SECTION V MULTILEVEL INPUT/OUTPUT IN THE CONTEXT OF LOGICAL SECURITY ENFORCEMENT	25
INTRODUCTION	25
THE CONCEPT OF MULTILEVEL INPUT/OUTPUT	25
THE MISSION OF LOGICAL SECURITY ENFORCE- MENT	28
IMPLICATIONS OF LOGICAL SECURITY FOR INPUT/OUTPUT	32

TABLE OF CONTENTS (Concluded)

		<u>Page</u>
SECTION VI	A GENERAL SECURITY MARKING POLICY IN THE ABSENCE OF MULTILEVEL I/O	35
	INTRODUCTION	35
	ESTABLISHMENT OF I/O TECHNIQUE GROUPS	35
	OUTPUT MARKING POLICY	38
	Group 1A	38
	Group 1B	39
	Group 2A	40
	Group 2B	41
	Group 3A	41
	Group 3B	42
	INPUT VERIFICATION POLICY	42
	Groups 1A, 2A, 3A	42
	Group 1B	43
SECTION VII	A GENERAL SECURITY MARKING POLICY IN THE PRESENCE OF MULTILEVEL I/O	45
	INTRODUCTION	45
	OUTPUT MARKING POLICY	46
	Group 1A	46
	Group 1B	47
	Group 2A	48
	Group 3A	50
	Group 3B	51
	INPUT VERIFICATION POLICY	51
	Groups 1A, 2A, 3A	51
	Group 1B	52
	Group 3B	52
SECTION VIII	THE IMPLEMENTATION OF TRUSTWORTHY LABELS	53
	INTRODUCTION	53
	CORRECT ISSUANCE OF LABELS	53
	PRECLUDING LABEL SPOOFING	54
	PRECLUDING LABEL ALTERATION	56
SECTION IX	SUMMARY	58
REFERENCES		59

LIST OF ILLUSTRATIONS

<u>Figure Number</u>		<u>Page</u>
1	The Interface Between Two Distinct Security Systems	7
2	Example of a "Tentative SECRET" cover sheet	14
3	Example of an incoming message (from AFM 10-2)	16
4	AFDSC Work Request/Receipt, AFHQ Form 34	23
5	The two threats with which logical security enforcement must deal	29
6	Three techniques for prevention of write-down	31
7	Input/output options of a Secret-cleared process	33

LIST OF TABLES

<u>Table Number</u>		<u>Page</u>
1	Definition of the groups of I/O methods, with examples	37

SECTION I

INTRODUCTION

The United States Air Force is currently making extensive use of computer systems to process classified information. A survey of ten Air Force commands has identified several existing systems, a wide range of anticipated future requirements, and a number of data security problems that are hindering the expansion of computer capacity to handle these perceived requirements. [1] Several projects at the MITRE Corporation are working towards short-term and long-term solutions to the computer security problems. One result of this work has been the development of the concept of a general security marking policy for classified computer input/output material. This project was performed under Sponsorship of the Electronic Systems Division.

Security markings are indications that are placed directly on, attached to, or included with classified material of any form. These markings serve "to inform and to warn the holder of the classification of the information involved, the degree of protection against unauthorized disclosure which is required for that particular level of classification, and to facilitate downgrading and declassification actions." [8, para. 4-101] The criteria used in applying and interpreting security indications constitute a marking policy.

A security marking policy, as the term will be used in this paper, has two objectives. The first is to fulfill the basic marking requirements of the security regulations for all types of computer I/O material on a consistent, generalized basis. This objective is important because the marking requirements for certain types of input/output techniques are vague and sketchy, and the requirements for a few other methods are non-existent. Establishing a general basis for setting marking standards would promote uniformity from installation to installation, and would facilitate the formulation of such standards for new I/O technologies as they become available for processing classified material. Also, the development of secure computer networks would be greatly aided by uniform standards for the labeling of data messages sent over communications lines.

The second objective of a security marking policy is to ensure that the security attributes of classified data items are accurately maintained on every I/O transfer. This objective recognizes that

the secure computer installation of the near future will employ two independent security enforcement systems. There will be a logical security enforcement mechanism that controls the access of processes (i.e., active programs) to data files and I/O devices, and there will be a physical security system that controls the access of personnel to classified documents and hardware. The logical enforcement mechanism will be one that depends upon internal machine representations of security labels to maintain security, while the physical enforcement system similarly depends upon document markings. The two systems will meet at the input/output interface as illustrated in Figure 1. Note that the security attributes of the two data items shown have been switched in the transfer across the I/O interface, all without violating either the logical or the physical security system. This is but one example of the ways in which the label or markings of a data item could be altered while that item is passing through the I/O interface. The alteration might occur as a result of a program bug, an operator error, or a malicious attempt to compromise classified information. If such an alteration were to take place, neither security enforcement system would necessarily detect it, since each one would observe only one side of the transfer. Therefore, the security marking policy must assume responsibility for establishing requirements and procedures that will prevent undetected security attribute changes from taking place.

One of the important aspects of the general marking policy is the development of a labeling policy. Labeling refers to the generation of security markings by the computer system itself during the course of an output operation. The labeling might be handled by the central processor, or it might be done by a separate secure communications processor which performs the bulk of the I/O processing for a large central computer. This application of security labels by a secure computer system is often considered a most significant feature of its operation, but wise decisions concerning the advisability and the scope of computer labeling can best be made in the context of an overall security marking policy.

This paper will consider the reasons behind the need for a general marking policy, review the guidance provided by existing regulations, and briefly outline the marking policies currently associated with three specific existing systems. It will then define the basic concept of multilevel I/O, and examine the impact of logical security enforcement on the usefulness of that concept. After a set of groups of input/output techniques have been established to help meet the first policy objective, two distinct security marking policies will be presented, one for unilevel I/O and one for multilevel. Finally, the paper will provide an overview of the considerations involved in the implementation of trustworthy computer labeling.




Figure 1. THE INTERFACE BETWEEN TWO DISTINCT SECURITY SYSTEMS

SECTION II

THE NEED FOR A GENERAL POLICY

Typical contemporary computer installations that handle classified information do not make use of a logical security enforcement system, in the sense of Figure 1. Rather, the machines run in a physically secure environment and process only one classification of data at a time. There is no risk of information being compromised (i.e., revealed to individuals not cleared to the appropriate classification level) within the computer, since all data is deemed to be at the same level. Whenever the security level of operations is changed, the system is sanitized (i.e., all demountable media classified at the old level are removed and all permanent storage is manually cleared or disconnected). The types of input/output media that are used (i.e., card decks, print-outs, magnetic tapes) have reasonably well defined security marking requirements. The computer system itself has no responsibilities for verifying or applying security labels; all marking responsibility rests with the operations and/or security personnel, and there are usually no doubts about the proper classification of any item. In those cases where doubts about classification arise, due to system error or other circumstance, the items are tentatively classified at the highest level for which the installation is cleared, and are given final classifications only after a painstaking review by authorized personnel. Under these circumstances, existing regulations suffice, and no further general security marking policy is necessary.

While this set of circumstances poses no unsolved information security problems, it does represent a very restrictive and inefficient utilization of computational resources. In order to satisfy continually increasing data processing requirements while keeping costs in line, it has become imperative to institute more flexible modes of computer operation. These new modes include the running of programs at several different classification levels concurrently on the same machine, the providing of on-line interactive service to users cleared to a variety of security levels, possibly including uncleared users who may pose security threats to the system, and the establishing of communications links between different military computer systems to form computer networks. As each of these improvements in computer utilization efficiency is implemented, a new problem will repeatedly arise: it will no longer be possible for the installation personnel to maintain data security without the active assistance of the computer system itself. At present, several efforts, directed at providing logical security enforcement for computers, are underway or have been proposed. [10]

The impact of this new situation on the security marking of classified computer input/output material will be dramatic. Relatively new input/output methods for which no firm security marking requirements exist (e.g., interactive dialog listings, CRT displays, audio I/O) will become commonplace, if not predominant. Also, the computer system itself will of necessity become one trusted source of information concerning the security attributes of active programs and their I/O requirements; it will indeed be the only trusted source of such information for work done on behalf of remote users. At the same time, it would be impractical to require a tape drive to sense tape reel colors, or to expect a card punch to stamp labels on the sides of card decks. Thus, the human operators and the computer system will have to share responsibility for verifying and applying security markings. Finally, it must be anticipated that new I/O methods will be devised and used for handling classified data, at which time new marking requirements, reasonably consistent with those for existing methods, will be needed.

This new and complex set of conditions brings about the need for a general security marking policy. Such a policy must establish a generalized basis for developing marking requirements appropriate to a wide range of input/output techniques. The policy must also set guidelines which will ensure accurate maintenance of security attributes on transfers between the physical security system and an assortment of logical security enforcement mechanisms. The application of this kind of general marking policy would promote development of operating procedures and standards that would be uniform from installation to installation and consistent from device to device. In its absence, the ad hoc development of confusing and incompatible standards would decrease the ability of different installations to share resources and increase unnecessarily the risk of compromising classified material processed by computers.

SECTION III

GUIDANCE PROVIDED BY EXISTING REGULATIONS

INTRODUCTION

A general security marking policy does not spring up out of a vacuum. In those areas where it overlaps existing regulations, it should agree with them. Elsewhere, it should be no more than a logical extension of current policies and procedures into new territory. It is, therefore, very important to review relevant regulations, both to determine what they say about the areas that they cover and to abstract general principles that can serve as a basis for those extensions which appear necessary. This section will be devoted to that review.

INFORMATION SECURITY PROGRAM REGULATION, DoD 5200.1-R and AFR 205-1

The regulation governing the classification, downgrading, declassification and safeguarding of classified information is DoD ISPR (Information Security Program Regulation) 5200.1-R [8]. The regulation amplifying those policies for use within the Air Force and providing procedural details where appropriate is AFR 205-1 [7]. Together, these regulations form the foundation for all Air Force policies and procedures regarding classified information. They share the same chapter and paragraph numbering system, and are to be considered as one unified document. Chapter IV, "MARKING", will be the subject of the following discussion.

Section 1 deals with the general provisions concerning marking of classified information. It states that "information determined to require classification protection against unauthorized disclosure... shall be so designated, generally in the form of physical marking." Every classified document is to show on its face its overall classification, whether it is subject to or exempt from scheduled downgrading and declassification, its office of origin, the identity of its classifier, the date of its preparation and classification, and, if appropriate, which portions are classified, at what level, and which are not. Material other than documents is to show such information on itself or in related or accompanying documentation. In addition, wholly unclassified material is not to be marked "Unclassified" except to convey that the material has been considered for classification and determined not to require it.

The specific requirements for classification markings on documents are given in Section 2. In general, "the overall classification of a document... shall be conspicuously marked or stamped at the top and bottom on the outside of the front cover (if any), on the title page (if any), on the first page, on the back page, and on the outside of the back cover (if any). Each interior page of a document shall be conspicuously marked or stamped at the top and bottom with the highest classification of information appearing thereon, including the designation 'Unclassified' where appropriate." Furthermore, "the classification marking must be in letters larger than those used in the text of the document, except in the case of documents produced on rapid printing automated data processing equipment." In that particular case, whether computer-printed documents or electrically transmitted record messages are involved, "these classification markings may be applied by that system, provided that the markings so applied are made clearly distinguishable on the face of the document from the printed text."

Requirements for paragraph marking within documents are also given in this section. These markings are different in nature from page markings in that the latter can usually be associated with specific device control functions (e.g., advance page), but paragraph markings cannot be. Also, it is unclear exactly how paragraph marking requirements might apply to such forms of output as program listings, tabular output, and graphics. Finally, the current state-of-the-art in computer data protection mechanisms tends to provide access controls only for quanta of information that are too large to support an internal analogy to paragraphs. For all of these reasons, a detailed discussion of general marking policy with regard to paragraph markings appears to be unwarranted at the present time. It can be assumed for the moment that when a computer-generated document properly requires such markings, the intended recipient of that item will be considered responsible for them.

In Section 3, the requirements for classification markings on material other than documents are listed. Certain computer-related materials, and their associated regulations, are as follows:

Magnetic Recording Tape - "Recordings, sound or electronic, shall contain at the beginning and end a statement of the assigned classification which will provide adequate assurance that any listener or receiver will know that classified information of a specified level of classification is involved... On reel flanges, mark the highest classification ever recorded on that tape. In addition, affix to the reel a label showing the current contents of the tape and classification data."

Magnetic Drums, Discs, and Disc Packs - "When removed from the processing machine, each individual drum, disc, or disc pack which has not been declassified...(overwritten or erased by a magnet) is marked with the highest classification of material ever recorded on it, and bears a label showing current contents and classification data. If stored in a container, the container also shows the highest assigned classification."

Decks of Accounting Machine Cards - "A deck of classified accounting machine cards need not be marked individually but may be marked as one single classified document so long as they remain within the deck... An additional card shall be added, however, to identify the contents of the deck and the highest classification involved. Cards removed for separate processing or use, and not immediately returned to the deck after processing,... shall be marked individually..."

The remaining sections of the regulation deal with the form and text of downgrading and declassification markings, provisions for the re-marking of old material and the form and text of certain specific additional warning notices. The notices provided for in the final section include those for Restricted Data, Formerly Restricted Data, Critical Nuclear Weapon Design Information, Sensitive Intelligence Information, and other information which is to be furnished to persons outside the Executive Branch. Such special warning notices need appear only once on a classified document or item, either along with or instead of the downgrading and declassification notice.

Two other chapters of the ISPR should be noted with respect to marking policy. Chapter XI deals with foreign origin material, and specifies additional warning notices to be used on NATO, CENTO, and SEATO material. It also outlines procedures for placing English-language classification markings on items which bear only foreign-language markings. Chapter XII sets out overall policy for dealing with special access programs. No details concerning the marking of items included in such programs are given, and it must be assumed that each program is free to establish its own marking requirements.

MANAGEMENT OF DATA PROCESSING EQUIPMENT, AFM 171-9

The manual which details the administrative policies applicable to the management of automatic data processing equipment (ADPE) throughout the Air Force is AFM 171-9 [6]. Paragraph 8-10 outlines the information security policies which apply "to all Air Force ADP

systems except computers used in command and control, communications systems, and computers integral to weapons systems." The underlying philosophy behind these regulations is given in subparagraph a, entitled "General Information":

"The rapid growth in the use of ADPE to process classified information has complicated the problem of safeguarding such information when processed or stored by the ADP system...Software and hardware, as presently operating, were not designed specifically with the objective of safeguarding defense classified information. This often creates a conflict between the desires of the user to operate the ADP system, using all of its capabilities, and the mandatory requirements for safeguarding classified information. Obviously, an accommodation must be effected which will neither degrade the efficiency of the ADP system unnecessarily nor subject the classified information to the unnecessary risk of compromise... Providing the necessary security for an item of classified information in an automated system is considerably more costly than protecting the same item of information in a manual system... This situation can be aggravated by overly stringent application of security procedures and overclassification. Therefore, it is incumbent upon all data processing installation (DPI) management and operating personnel to constantly seek improved methods to ensure security, consistent with basic security policy established in AFR 205-1 and in this paragraph."

Remaining sub-paragraphs deal with, among other things, degaussing/erasure/overwriting of classified information, security considerations in serial job processing ADP systems, and security considerations in resource-shared ADP systems. The marking of classified I/O material is dealt with only under resource-shared systems, and even there the only provision established is one concerning cover sheets appropriate to print-outs from systems which schedule unclassified, Confidential, and Secret applications together. Such cover sheets instruct the recipient to treat the print-out as a Secret document until appropriate inspection leads to a final classification. One example of a cover sheet conforming to this requirement is AFHQ Form 0-421, shown as Figure 2.

THIS COVER SHEET IS UNCLASSIFIED

SECRET

THE FOLLOWING STATEMENT APPLIES TO THE ATTACHED DOCUMENT
AND IS STATED IN ACCORDANCE WITH PARAGRAPH 8-10 D (5),
AFM 171-9 (C20), 1 JULY 1970.

WARNING

This document may contain CLASSIFIED INFORMATION, until it
has been reviewed it will be controlled, accounted for, stored, and
transmitted as SECRET. When the requester has assured himself
that it contains no data not requested the tentative classification
will be affirmed SECRET or regraded to CONFIDENTIAL, or de-
classified UNCLASSIFIED. If this document contains data not re-
quested, the entire document will be returned to the Air Force Data
Services Center (AFDSC). Room 1D 1080, The Pentagon, Wash.,
D.C. 20330, for review/destruction. All control / accountability
procedures consistent with tentative SECRET CLASSIFICATION
will be maintained.

SECRET

THIS COVER SHEET IS UNCLASSIFIED
"NOTE" SEE REVERSE SIDE

AFHQ 0-421

Figure 2. Example of a "Tentative SECRET" Cover Sheet

Implicit in this regulation is the notion that no effective internal computer security mechanisms exist, and therefore, when operational needs dictate multilevel processing, the computer installation itself cannot be made responsible for placing final security markings on I/O material. So long as that initial assumption is valid, the conclusion is equally valid. However, adequate internal mechanisms are currently under development and will soon be available. Therefore, a general security marking policy which will be useful in the near future must insist that every DPI assume full responsibility for applying final markings to all output products and verifying the security markings of all input material.

AIR FORCE MESSAGE MANAGEMENT PROGRAM, AFM 10-2

The manual which establishes the administrative policies, procedures, and standards applicable to the management of record messages throughout the Air Force is AFM 10-2 [5]. Although this manual does not deal directly with the security marking of computer input/output, it does address the requirements for classified record messages printed by automated equipment, as well as standards for magnetic tape and data pattern (card deck) messages. Thus, it might provide useful guidance.

The section dealing with magnetic tape and card messages provides that such messages must be accompanied either by AUTODIN header and End of Transmission (EOT) format cards or by a completed DD Form 1392, "Data Message Form", which would enable the telecommunications center to prepare such cards. In either case, the cards are used to transmit all message control information, including security classification and any other special security control information.

Attachment 9 to this regulation, reproduced here as Figure 3, shows an example of an incoming message printed by automated equipment. Of particular interest are the "SECRET" marks at the top and bottom. They demonstrate one accepted method of differentiating between classification label and text, namely, surrounding the label by asterisks.

TECHNIQUES AND PROCEDURES FOR IMPLEMENTING, DEACTIVATING, TESTING, AND EVALUATING - SECURE RESOURCE-SHARING ADP SYSTEMS, DoD 5200.28-M

The manual implementing DoD Directives and Instructions, and establishing uniform guidelines for techniques and procedures to be used when implementing, deactivating, testing, or evaluating secure resource-sharing ADP systems, is DoD 5200.28-M [9]. The manual in its present form presents only the most broad and basic guidance, except in those areas where operational experience provides an

THIS SPECIMEN IS UNCLASSIFIED

***** *S E C R E T* *****	RTTU JAW RUEFQA 1710 3130849 SSSS	Four letter classification code as entered on DD Form 173 originator.
ZNY SSSS	R 090845Z NOV 69	This is the UNICOM symbol, the originator and office symbol, identifies the message. The first two digits represent the date, the next four digits the time (Zulu time). When referring to this message you would state: CSAF (AFMSPA) 090845Z Nov 69. When replying to this message you would state: four AFMSPA 090845Z Nov 69.
FM CSAF WASH DC	TO RUHLQH/CINCPAC CANP H SMITH HAWAII/DA	Action to these addressees.
RUHLKM/CINCPACAF HICKAM AFB HAWAII/DA/DE	RUHUAAX/SAF FUCHU AS JAPAN/DAS/CEN	Information to these addressees.
INFO RUMABA/13AF CLARK AB PHIL/DAS/OSO	RUCCEN/VRAMA ROHINS AFB GA	Office symbol of originating agency.
BT	S E C R E T LIMDIS DA	
FOR DIR ADMIN. SUBJ: SAMPLE MESSAGE (U).	REF: A. MY AFMSPA 071830Z NOV 69	
B. JCS (J6) LETTER, 8 NOV 69, SUBJ: MESSAGE REFERENCE (U)	C. ADC (ADCCR) 090800Z NOV 69	
THIS IS A SAMPLE OF AN INCOMING MESSAGE AS RECEIVED ON A TELETYPE MACHINE. YOU WILL BE PRIMARILY INTERESTED IN THE EXPLANATION SHOWN. CP-A	BT	Group coding will be reflected here.
***** *S E C R E T* *****		Sample classification marking applied by an automated comm system. Authorized in lieu of stamps, or other markings. See AFR 205-1. Messages received unmarked should have proper security classification markings entered per AFR 205-1. Local instructions will dictate the office responsible for marking.

Figure 3. Example of an Incoming Message (from AFM 10-2)

understanding of detailed requirements (e.g., magnetic tape erase procedures). It is geared toward the future expansion, augmentation, and revision that will occur as experience with secure resource-sharing ADP systems begins to develop.

While little specific guidance in the area of an I/O security marking policy is given by the manual in its present form, certain requirements are listed. Section IV deals with those hardware and software features deemed essential to provide protection for classified material, and Part 3 of that section is devoted to software features. One of the listed features is "Security Labels", for which the given requirement is as follows: "All classified material accessible by or within the ADP System shall be identified as to its security classification and access of dissemination limitations, and all output of the ADP System shall be appropriately marked."

Section V adds the requirements for an audit log or file. This log may be manual, automatic, or some combination of the two. It "shall be maintained as a history of the use of the ADP System to permit a regular security review of system activity." Examples of specific transactions to be recorded include "logins, production of accountable classified outputs, and creation of new classified files."

ABSTRACTION OF GENERAL PRINCIPLES

This has surely not been an exhaustive review of security-related or computer-related DoD or Air Force regulations. However, other potential sources of guidance tend to repeat the material that has been reviewed or to be even more vague and general than those citations that have been given. Therefore, at this point, a list of general principles may be drawn up to serve as the basis for a general I/O security marking policy. The list, based entirely upon the material included in this section, is as follows:

1. Every discrete, separately handled item of classified material requires indications of:
 - a) overall classification level
 - b) declassification schedule or exemption category
 - c) source of classification authority
 - d) date of production/classification
 - e) additional warning notices, as appropriate

2. Marking of the classification level of components within a separately handled item is sometimes required (e.g., document pages) and sometimes not required (e.g., punched cards within a deck).
3. If the classified contents of an item are written text, the classification level marking must be distinguishable from the text. One accepted method of achieving distinguishability on a high-speed printer is to surround the classification marking with a border of asterisks.
4. If the classified contents of an item are other than written text, written marking is nevertheless always required. However, an additional indication of classification level incorporated into the classified contents is sometimes required (e.g., voice recordings) and sometimes not required (e.g., magnetic discs used for digital data storage).
5. An audit log will record every instance of the production of accountable output items or the creation of classified files from accountable input items.

SECTION IV

EXISTING POLICIES FOR SPECIFIC SYSTEMS

INTRODUCTION

In addition to the guidance provided by current regulations, the experience of those currently responsible for the processing of classified material on computers can be of great value in the formulation of a general marking policy. Of course, no contemporary system can be expected to provide examples of how to ensure the accurate transfer of security attributes across an I/O interface. Nevertheless, local standards for the marking of classified data processing material have been established for several installations, and these policies should be reviewed before dealing with the first marking policy objective, that of creating a generalized basis for setting marking requirements. At best, these local standards will yield worthwhile insights into how certain types of input/output media should be handled. At worst, they may at least further demonstrate the need for a general policy. Three examples of existing policies will be reviewed in this section.

MITRE CORPORATION COMPUTER FACILITY

The MITRE main computer facility consists of an IBM 370/158 system and associated peripheral equipment. During normal working hours, the facility is available for batch or time-shared processing of unclassified material only. At other times, it may be used for the processing of Confidential or Secret information on a dedicated basis. The procedures applicable to such classified processing are outlined in Volume V of MITRE Security Procedures, entitled "Computer Operations."

Paragraph 106 of this volume is entitled "Marking of Classified Materials." It is brief and concise enough to quote in its entirety:

- "Card decks are marked as a unit or as individual cards. When a deck is handled as a unit, the top side of the deck (as viewed when the deck is in a tray) shall be conspicuously stamped with the highest overall classification. There will be a security header card for each deck marked with the following information: classification, name and address of the facility, subject or title, date, downgrading notice, espionage notation, document control number and number of cards.

- "Classified printouts are always made on preprinted paper stock (request SYSOUT = S for SECRET; SYSOUT = C for CONFIDENTIAL). The name and address of the facility, subject or title, date, downgrading notice, espionage notation, copy number, and document control number will appear on the first page of each printout.
- "Tapes and disk packs must have all required security markings listed above placed conspicuously on each reel or pack. Computer Center US Series tapes have distinctive red reels for ease of identification.
- "Material will not be accepted by the Computer Center unless properly marked as described above."

The reference to US Series tapes is explained by a section of Paragraph 104: "Materials logged into the Computer Center are normally stored and used there until the job is completed, in any case, not more than thirty days. For jobs requiring a longer period of time, the Computer Center has set aside a group of labeled tapes (US Series) for permanent use in the facility. The US Series tapes, identified by the red reels, are to be stored separately from other tapes. These tapes fall into two categories: 'save tapes' and 'scratch tapes'. Save tapes have classified information on them. Scratch tapes are unclassified tapes reserved for classified runs."

The MITRE installation is typical of contemporary facilities that perform classified electronic data processing on a dedicated system basis. Because of the strictly controlled nature of the processing, cleared operations personnel can easily assume full responsibility for all application and verification of required media markings. No internal security system is needed or used. However, dedicating an entire facility to a single job is very expensive and inefficient. The remaining two example systems operate in more complicated fashions.

WWMCCS GCOS-III SYSTEM

The World Wide Military Command and Control System (WWMCCS) is a vast network of computer installations, surveillance sensors, and communications links serving the National Command Authority, the Joint Chiefs of Staff, and other U.S. military commanders. To meet a requirement in the WWMCCS ADP contract, a security package for the H6000 General Comprehensive Operating System (GCOS) was developed by Honeywell Information Systems, Inc. (HIS). The basic features offered by this security package are described in an Operating

System Technical Bulletin issued by the Joint Technical Support Activity [13] and a Series 6000 Software Manual published by HIS [12]. It must be emphasized that these security features are not intended to enforce the rules governing protection of classified information. Rather, they are intended "to give the people ultimately responsible as useful and as versatile a set of tools as could be devised to enable the user to manage classified data...if these [tools] cannot force a user to manage classified data wisely they can at least force his attention to the fact that he is dealing with classified data and force him to take certain steps to provide for the protection of that data." [13, p.1]

The security marking policy of the WWMCCS GCOS Security Features deals exclusively with system markings on batch and terminal printed output. In the batch case, it allows each facility to define the labels available for printing at the top and bottom of every page of printed output, and it asks that each user specify which label is to go on his or her output. The classification and category set specified for the job are printed on the output only as a default condition, if no other classification code is included on the command card generating the output. In addition, provision is made for the installation to define and the user to select a code specifying "DO NOT MARK", i.e., no labels to be printed on the output. No other forms of batch output are dealt with, verification of input markings is not considered, and no mention is made of additional required markings.

In a terminal mode, the log-on conversation requires the user's selection of the installation defined security labels to be printed on the output. In response to the question, "CLASSIFICATION OF YOUR OUTPUT?", the user may give any valid classification code, including those representing "DEFAULT" and "DO NOT MARK". If files are to be created, a separate question, "CLASSIFICATION OF FILES YOU WILL CREATE?", is asked. The response may be any legal code except "DEFAULT" and "DO NOT MARK"; in particular, the code need not be that given for output marking. Again in this case, it seems that many factors have never been considered.

There appears to be little to learn from the WWMCCS security marking policy. The output labels are, of course, untrustworthy since the operating system is uncertifiable. Even if they could reliably reflect the security level of the data they accompany, their timing in terminal mode is uncertain without a clear definition of "page tops and bottoms" in interactive dialogue. Finally, even if the labeling scheme for printed output were fully acceptable, and if the operator who removes the printed material from the printer were made fully responsible for all additional markings, then there

are still no marking standards for any other form of input/output. In fact, the labels that are provided for are not standardized from installation to installation. The WWMCCS GCOS-III System is in need of a general security marking policy; it appears to have very little to contribute toward one.

UNITED STATES AIR FORCE DATA SERVICES CENTER

The Air Force Data Services Center (AFDSC) is an installation located in the Pentagon which provides general purpose data processing services to the Air Staff and the Office of the Secretary of Defense. This facility is cleared to handle material classified up to and including TOP SECRET, and its computer resources include three HIS G635's, an HIS 6060 WWMCCS machine, and HIS 6180 MULTICS machine, an IBM 360/75, and an RCA SPECTRA 70 AUTODIN processor. The input media used by the Center include punched cards, magnetic tapes, disk packs, and keyboard entry, while the output media include the first three input items, plus plotter output, typewriter print-out, CRT images, and line printer output.

Procedures exist which provide physical security and personnel entry/exit control for the main computer area, the tape library area, the bindery area, the keypunch area, and an assortment of remote sites, all located within the Pentagon complex. Additional remote sites, equally secured, are located outside of the Pentagon. The remote sites contain interactive and remote batch entry terminals, and communications security is provided for the lines connecting the terminals and the main computer area. The provisions for physical security and other procedures dealing with the operation of the United States Air Force Data Services Center are documented in a Security Procedures Manual.

Since the AFDSC handles multiple levels of classified data, the procedures include provisions which amount to a de facto marking policy for classified computer I/O material. For example, AFHQ Form 34, illustrated in Figure 4, is used to help keep track of the classification of input material and to indicate the expected classification of output. However, the general operational philosophy reflected in the Security Procedures Manual states that the computer operations personnel cannot ascertain the exact security status of any particular item of output material. This situation arises because throughput requirements dictate that machines concurrently process information at different classification levels, but adequate logical security enforcement mechanisms are not yet available. As a result, the Center must rely heavily upon the provisions of AFM 171-9, cited in the preceding section,

DATA SERVICES CENTER G-635 WORK REQUEST / RECEIPT										SNUMB				
NAME		OFFICE	PHONE	HOME PHONE	DATE		SYSTEM ☐ A ☐ B			ABORT CODES				
SNUMB	ACCT NUMBER	DSD	TYPE TIME	INSTRUCTIONS ON BACK		RECYCLE (OF) (AS)								
TIMER EXTENSION ☐ YES ☐ NO		SYSTEM ☐ A ☐ B ☐ EITHER		INPUT CLASS		TS		S		C	FOUO	U		
RECYCLING PERMISSABLE ☐ YES ☐ NO		☐ READ PRMFL		☐ CREATE OR UPDATE PRMFL		☐ NO PRMFL USE		OUTPUT CLASS		TS	S	C	FOUO	U
MAX RESOURCES REQ		NUMBER OF ACTIVITIES		NUMBER OF BMC PUNCHES		NUMBER OF BMC PRINTS		NUMBER OF OTHER PUNCHES						
TAPE DRIVES _____ CORE _____		CPU TIME _____ TEMP _____		MASS _____		RMVBL DISC DRIVES _____		SYSOUT LINES _____						
REEL/PACK NUMBER		SIGNATURE OF APPROVING OFFICIAL		SIGNATURE OF RECEIVER										
AFHQ FORM 34, OEC. 71, REPLACES AFHQ FORMS 34, 35, 36 DTD JAN 70, WHICH ARE OBSOLETE.														

Figure 4. AFDSC Work Request/Receipt, AFHQ Form 34

that allow installations such as the AFDSC to apply tentative classifications to computer-generated output material.

The Air Force Data Services Center is essentially a benign environment in the sense that uncleared users are not authorized to utilize machines performing classified processing. Its computers must in general go through tedious sanitization procedures in order to change security levels, though concurrent processing of multiple levels without logical security enforcement is often necessary. Its marking policy reflects this status. However, the AFDSC is presently engaged in revising its security procedures in an attempt to accommodate a broader range of users on the multilevel secure systems of the near future. The remainder of the paper will deal with some of the issues which this facility, and others like it, are now beginning to face.

SECTION V

MULTILEVEL INPUT/OUTPUT IN THE CONTEXT OF LOGICAL SECURITY ENFORCEMENT

INTRODUCTION

With a review of relevant regulations and existing policies completed, the substance of a general security marking policy may now be considered. This section, however, will not lay out policy details, but rather will introduce a fundamental concept, that of multilevel input/output. The distinction between unilevel I/O and multilevel I/O is so basic to the issues addressed by security marking policy that two separate policies, one for each variety of I/O, will be developed. Therefore, it is essential to define the notion of multilevel input/output and examine the applicability of this concept in an environment of logical security enforcement before the policies themselves can be discussed.

THE CONCEPT OF MULTILEVEL INPUT/OUTPUT

The notion of multiple levels of sensitivity and protection is common to most systems that seek to restrict the dissemination of information. Multiple levels arise because it would be overwhelmingly expensive to protect to the maximal extent all material requiring any protection, and it would be decidedly impractical to clear to the highest level everyone who requires access to any protected material. Therefore, a set of levels is defined that permits a practical degree of protection to be provided for information, corresponding to its degree of sensitivity.

The United States military information security system defines levels of sensitivity by using two variables. The first is classification, a hierarchical set including (but, under special circumstances, not limited to) four levels: Unclassified (i.e., requiring no protection), Confidential, Secret, and Top Secret (i.e., requiring maximal protection). By hierarchical, it is meant that the four classifications are ordered, and clearance to one level implies clearance to all lower levels. The second variable is access category, a non-hierarchical set orthogonal to classification. Access categories are not fixed in number; new ones may be created and old ones may be terminated. At present, they include material protected at the request of foreign powers (e.g., pact organizations such as NATO and SEATO, as well as the governments of individual foreign nations), material protected under the authority of non-military agencies (e.g., Restricted Data under the Atomic Energy

Commission, CRYPTO material under the National Security Agency, etc.), and material associated with specific restricted access programs entirely within the military.

Standards of physical protection are defined for material at each level of sensitivity. For items at specific classifications and in no special access categories, the levels of physical protection correspond exactly to the classification levels. For an item at a particular classification and in certain categories, the level of physical protection mandated may be somewhat more elaborate than that for no-category items of the same classification, but still not adequate for the protection of material at the next higher classification. An item in certain other categories may require more elaborate physical protection than no-category Top Secret material, irrespective of the item's actual classification. In any case, the United States military security system can be spoken of as assigning every item to one of a single ordered set of physical protection levels, with those levels defined by the two security variables, classification and category.

Of course, it is not a security violation if material at a particular sensitivity level is physically protected to a level higher than the mandated one. However, making this sort of overprotection a regular practice is deemed undesirable, primarily for two reasons. First of all, higher physical protection levels are always more expensive to create and maintain. Therefore, the overprotection of significant amounts of material generally represents a serious waste of limited funds and resources. Secondly, continual overprotection can often lead to carelessness on the part of personnel responsible for maintaining security. People tend to be somewhat casual about handling material as if it were exceedingly sensitive when they know that it usually is not.

These considerations have a direct bearing on the operation of data processing installations that handle classified information. The physical protection accorded the computer itself must, of course, correspond to the highest sensitivity level of information that can ever be processed by the installation. Some input/output devices, however, may be restricted to processing information only at levels well below the installation's overall clearance, and a correspondingly lower level of physical protection should be provided for these devices. Indeed, the protection provided for each individual item of I/O material should be only as stringent as is dictated by the sensitivity of the item's contained information. This philosophy implies that the installation personnel should always be cognizant of the level of physical protection which is appropriate to each item of I/O material that they handle.

Contemporary computer systems generally do not include logical security enforcement mechanisms. As a result, they must process

information at only one sensitivity level at a time, in order to avert the possibility of unauthorized persons obtaining classified information by essentially instructing the computer to give it to them. Systems which operate under these conditions may be called unilevel systems. One characteristic of a unilevel system is that it must undergo sanitization when the sensitivity level of its computational load changes. The nature of sanitization has been outlined earlier (see Section II). In this environment, it is an easy matter for installation personnel to know the sensitivity level of all I/O material, since each I/O device can at any instant only be processing material at the level of the system itself. If proper procedures have been followed, all other material has been removed from the machine room or locked in appropriate storage.

Newer systems will incorporate logical security enforcement mechanisms. The actual amount of increased operating flexibility obtained will depend on the comprehensiveness and reliability of the mechanism employed. In any case, such systems will be multilevel systems, authorized to concurrently perform computation at more than one sensitivity level. Each input/output device can, in this new environment, operate in one of two modes. One mode involves making the system appear, from the point of view of the I/O device, as unilevel rather than multilevel. The device itself processes information requiring only one level of physical protection. To change the level of material handled by the I/O device, a security reconfiguration must be performed. This procedure is characterized by the conscious intervention of the system security officer, who must alter the data base of the system's logical security enforcement mechanism in order to effect the reconfiguration. This mode of I/O device operation is referred to as unilevel I/O in a multilevel system.

Other I/O devices attached to the system may operate in a more sophisticated fashion. The entry for these devices in the security mechanism's data base will not be the single level of information that they will handle, but will rather be a device operating clearance, the maximum level which the device may currently see. (Of course, the device will also have a permanent ceiling clearance, based upon the physical protection provided in the area where the device resides). A security reconfiguration will be required in order to change the operating clearance of one of these devices. However, at any instant, these input/output devices will be authorized to process either information requiring physical protection at their operating clearance level, or information requiring any lesser degree of protection. Within the authorized range, the actual level being handled at any moment by a device will be decided on the basis of efficiency considerations rather than security ones. In particular, no human intervention will be required for level changes that remain within the range of device's operating clearance. This type of operation is known as multilevel I/O.

At this point, a word about the notion of security enforcement by a computer system should be added. People understand the need for security and the consequences, both to the nation and to themselves, of compromising security; machines do not, and cannot be expected to in the future. However, the technology to produce provably correct computer programs exists, and the methodology for applying this technology to the creation of certifiable logical security enforcement mechanisms is being developed. [10] These advances will permit appropriate officials to guarantee, on their personal authority, that various security-related functions will be performed correctly by a computer system. In this sense, software can be certified, and certified software may be said to be "responsible" for security enforcement, just as an approved safe may be said to be "responsible" for the physical protection of classified documents.

Unilevel I/O, whether in unilevel or multilevel systems, requires no further discussion before its marking policy can be dealt with. The same is not true of multilevel I/O. This latter mode of operation cannot exist without the support of a logical security enforcement mechanism. Yet, paradoxically, the very task carried out by logical security limits the extent to which multilevel I/O can take place. The remainder of this section will be devoted to explaining the limitations on multilevel I/O. The mission of logical security enforcement will be described, and then the implications for input/output of carrying out that mission will be pointed out.

THE MISSION OF LOGICAL SECURITY ENFORCEMENT

In multilevel systems, there are two ways in which information might be revealed to unauthorized persons that physical security arrangements are not equipped to avert or even detect. These techniques for achieving security compromise are referred to as read-up and write-down. [11] The prevention of read-up and write-down constitutes the primary mission of logical security enforcement. (Logical security also involves secondary issues, such as the maintaining of accountability records, but the actual prevention of security compromise consists entirely of blocking the two named techniques). These two basic threats are illustrated in Figure 5.

Read-up involves a process, acting on behalf of a specific individual, reading information resident in the system which the individual himself is not cleared to read. This security threat is conceptually very straightforward. Thwarting it consists of maintaining records of the classification and category of all passive system resources, as well as records of the authorization of each process, and preventing processes from reading resources that they are not authorized to read.




Figure 5. THE TWO THREATS WITH WHICH LOGICAL SECURITY ENFORCEMENT MUST DEAL


While the mechanisms that claim to perform this task may vary widely, they all take essentially this same conceptual approach.

The other threat, that of write-down, is a somewhat more complex problem. Its scenario consists of a sequence of actions. First, a process reads sensitive information from an area to which it has legitimate access, and then it writes that data into some other area which the system considers less protected or unprotected. Finally, a second process, authorized to access only the latter area, reads the sensitive information that has been placed there. In this way, the security of a multilevel system can be compromised without any read-up violations taking place. The threat of write-down can be handled in one of three basic fashions, as illustrated in Figure 6.

One method is the *-property technique, identified by Bell and LaPadula [2]. Under this technique, the impact of each access request is assessed before the request is granted. Any request which would lead to a given process having read access to one resource and write access to a less protected resource concurrently, is denied. In this way, the chain of accesses implicit in the write-down threat can never be formed.

The high-water mark technique, used by Weissman in the ADEPT-50 time-sharing system [14], offers a different approach. It permits the write-down access chain to be formed, and it permits the write operation to take place. However, it then upgrades the written-into resource to the level of the read-from resource. Thus, the second process can no longer read the transferred data except by creating a read-up violation, and read-up is already controlled by the logical security enforcement mechanism.

Finally, one method for preventing write-down permits the access chain creation, the write operation, and the final read by the second process all to take place. However, this may occur only under two strict conditions: all software involved in the "first process" must be certified to ensure that no sensitive information will be transferred to an area less protected than appropriate, and all data in the first read-from system resource must be labeled to indicate the actual degree of sensitivity. This is the trustworthy software approach, suggested by Bell in a modification of the *-property [3]. Very few systems can be expected to utilize certified software only, so this last technique must usually be used in conjunction with one of the other two.



18-42,971

Figure 6. THE THREE TECHNIQUES FOR PREVENTION OF WRITE-DOWN

These, then, are the problems that logical security enforcement mechanisms must face, and the approaches available for dealing with them. At this point, it must be recalled that system resources include not only memory and other permanent on-line storage, but also the various input/output devices associated with the system. The notions of read-up and write-down prevention must be examined in terms of their impact on I/O.

IMPLICATIONS OF LOGICAL SECURITY FOR INPUT/OUTPUT

The ways in which logical security enforcement affects I/O in a multilevel system can best be understood by imagining a process that runs on behalf of a Secret-cleared user, and considering what its input/output options are. A small assortment of I/O devices can also be postulated: three line printers, three card readers, and three tape drives, one of each type cleared to Confidential, Secret, and Top Secret. Special access categories can, for this discussion, be ignored.

The line printers represent the general class of write-only devices. Their situation is pictured in Figure 7a. For the Secret-cleared process, producing output on the Confidential-cleared printer would constitute write-down, and would be prohibited (except in the trustworthy software case, when only Confidential print-outs would be produced). This is reasonable; the printer might well be physically protected only to the Confidential level and manned by operators with only Confidential clearances, making the printing of Secret material on the device improper. However, the logical security enforcement mechanism would do nothing to prevent the process from causing unclassified, Confidential, or Secret output through either of the other two printers. In addition, Top Secret activity on the system could result in Top Secret output through the printer cleared to that level. It is apparent that multilevel output is a distinct possibility through write-only devices.

In view of this situation, it might be assumed that multilevel input through a read-only device is likewise allowed. However, as Figure 7b illustrates, this is not the case. For the Secret-cleared process to obtain any input through the Top Secret card reader would be a read-up violation. For the process to accept input from the Confidential reader, it would have to be able to place that input into Confidential storage. However, unless the trustworthy software case applies, the process is prevented from performing write-down, and so it could not consummate the transfer of Confidential material into memory without over-classifying it. Therefore, in order to avoid large-scale over-classification, the Secret-cleared process is not permitted to accept input through the Confidential-cleared card reader.






Figure 7. INPUT/OUTPUT OPTIONS OF A SECRET-CLEARED PROCESS

Of course, the same applies to less-than-Secret material read in through the Secret-cleared device. This basic fact generalizes to processes running at other sensitivity levels, and leads to the principle that multilevel input does not take place even on multilevel systems, except under the conditions imposed by the trustworthy software technique.

The status of the tape drives, which represent read/write devices, should now be clear. As is shown in Figure 7c, access by the Secret-cleared process to the Top Secret-cleared drive constitutes a read-up violation, while having it use the Confidential-cleared device amounts to initiating write-down. Read/write access is available only to the device at the process's own level, and even the very data on the tape itself can only be Secret, since more sensitive material's presence would compromise security and less sensitive material would get overclassified. Of course, the usual provision must be made for the exceptional circumstances which permit write-down by trustworthy software. In any other case, though, read/write devices are not permitted to run in a multilevel mode.

The specific types of devices mentioned in this discussion were only intended to serve as illustrative examples. Actual policy for specific I/O device types will be discussed in the sections that follow. In particular, it should be noted that tape drives need not always act as read/write devices. Restricting a drive to using only tapes without fixed write-protect rings, for example, converts that drive into a read-only device.

The proper context has now been established for the presentation of the two general security marking policies, one for unilevel I/O and one for multilevel I/O. The unilevel policy was always the only one that could ever have applied to unilevel systems. Now, though, it appears that unilevel policy will have very wide application to multilevel systems, as well. Multilevel policy, in fact, will be useful only for write-only output devices, other I/O devices controlled exclusively by trustworthy software, and one other class of devices that now deserves mention.

This third possibility will prove vitally important in dealing with interactive terminals. It will be possible for trustworthy software to control a multilevel input device, set that device to a particular single level, and then turn it over to non-trustworthy software. The sensitivity level of the material handled by the device can only be changed when trustworthy software regains control. Nevertheless, the only type of policy appropriate to the device would be a multilevel one, even while the non-trustworthy software is in control. Further discussion of this case, as well as detailed consideration of the whole range of I/O cases, will be found in the exposition of the two general security marking policies.

SECTION VI

A GENERAL SECURITY MARKING POLICY IN THE ABSENCE OF MULTILEVEL I/O

INTRODUCTION

This section presents a general security marking policy specifically aimed at input/output devices that exclusively perform unilevel I/O. Before the devices addressed by this policy can change the sensitivity level of their traffic, they must be involved in either a system sanitization or a system security reconfiguration, depending on the type of system with which they are associated. This implies that operators of these devices will be notified of any level changes and will have ample opportunity to physically indicate such changes on the devices themselves. In this way, the second policy objective, that of maintaining accurate transference of security attributes, can readily be met.

To facilitate the establishment of a general marking policy, a set of groups of input/output techniques will be defined through the identification of two group dimensions. Then, technique examples and policy will be given for each group, first considering only output marking and later considering input verification.

ESTABLISHMENT OF I/O TECHNIQUE GROUPS

The initial task involved in meeting the first marking policy objective, that of creating a consistent and generalized basis for fulfilling the basic marking requirements, is the establishment of groups of input/output methods. These groups must be specific enough to distinguish between techniques with substantially different security marking requirements, and yet general enough to include all I/O methods that exist or are realistically expected to be developed. The creation of this set of groups is facilitated by the identification of two group dimensions, removability and legibility.

Removability refers to the potential for removing input/output material from the device through which that material has been entered into or produced by a computer system. In this dimension, there are two possibilities: I/O material may be removable or non-removable. Removable material has independent physical substance and may exist separately from any automatic data processing equipment. This type of medium can be compromised by removal to an unsecured area, and must be stored in a physically secure fashion until destruction.

Non-removable material, on the other hand, exists solely as the physical state of one or more component parts of input/output hardware. This material is not liable to compromise by physical removal (except by the removal of the I/O device itself, which may be considered unlikely), and it may be destroyed merely by setting the device to a neutral state. Since the latter type of I/O method inherently entails substantially less risk of compromise, its security marking policy may be significantly less stringent than the one for removable I/O material.

The second group dimension is legibility. This attribute refers to the degree of machine assistance normally required for a person to apprehend various forms of I/O information, and it includes three possibilities. Directly human-legible material may be interpreted through the direct application of human senses, without any recourse to artificial aids. Then, there is other material that can only be interpreted through the use of some device, but for which the required device does not perform any digital data processing. Finally, there is material that is only normally interpreted by a computer system or other digital information processing equipment. (The "normally" must be emphasized; there are people who are quite adept at reading the holes in punched cards, for example. The objective of distinguishing the third possibility is to highlight those items most likely to be involved in multilevel input). While all forms of classified material will require certain minimal directly human-legible security markings, general principle 4, identified in Section III, indicates that the legibility dimension may in some cases generate additional special marking requirements.

The two group dimensions can be used to establish a set of six groups of input/output methods (of which only five are meaningful, as will be seen). Table 1 indicates how these groups are specified, and applies an identifying label to each one. In subsequent subsections, the unilevel marking requirements for each group of methods will be discussed, and examples of the application of these requirements to specific I/O techniques will be given. The more obvious task of applying output markings will be considered first for the entire set of groups, followed by a review of the set with an eye toward the task of verifying input file attributes.

Before dealing with specific cases, one general point should be mentioned. General principle 5 of Section III requires that an audit log be maintained, and it would seem that this requirement is quite appropriate to unilevel resource-shared systems as well as to multilevel ones. However, the concept of auditing in secure computer systems has several connotations, and it is necessary to make clear exactly which one is involved in marking policy.

Table 1

Definition of the Groups of I/O Methods, with Examples

REMOVABILITY

		A	B
		Removable	Non-removable
L E G I B I L I T Y	1	Directly human-legible	
	2	Human-legible with non-digital machine aid	
	3	Digital machine legible, only	

Group 1A	Group 1B
(line printer, printed interactive dialog, X-Y plotter)	(CRT display, keyboard entry, voice I/O)
Group 2A	Group 2B
(microfilmed output, audio tape recording)	(meaningless intersection, no examples)
Group 3A	Group 3B
(punched cards, perforated paper tape, magnetic tape & discs)	(data messages sent via computer network links)

One notion of an audit log is an automatically maintained record of the incidents of creation of accountable items, analogous in function to the production log of a classified printing center. Another type of log is one that keeps a record of each incident of access by an individual to a classified item, analogous to a record of the opening and closing of a safe. Finally, the term is sometimes used to denote a record of incidents involving detected violations of security procedures or policies. The latter two connotations deal with issues beyond the scope of marking policy, but an audit log in the first sense is a major tool that permits an adequate policy to be implemented.

Throughout both general policies, then, it is assumed that the device operator has access to the "creation" audit log for the device in question, through the system security officer. It is further assumed that the information recorded in the log for each incident of creation includes as a minimum the classification/category of the item created, the user ID, the project ID, the declassification schedule/exemption code, and the source of classification authority, as well as the requisite date, time, and I/O device ID entries. This data should be sufficient to provide the operator with all of the information needed to apply those markings for which he or she is responsible.

OUTPUT MARKING POLICY

Group 1A

Directly human-legible, removable output material includes all items normally considered "documents" by the existing regulations. However, the marking requirements for documents are page-oriented, and not all printed computer output necessarily comes in pages. Therefore, it is necessary to divide this class into two sub-groups, paged and non-paged.

Paged printed output has the most well defined marking requirements of any I/O medium. Typical examples of this type of output are print-outs produced by line printers using continuous, fan-folded forms and drawings produced by hard copy X-Y plotters using single-sheet paper. One requirement applying to this case is that each page of output have its classification marked at the top and bottom. This marking is most readily applied for a unilevel device through the use of pre-printed forms. Indeed, the use of pre-marked paper permits the operator to respond to a device level change by simply changing the forms which are fed to that device. The other important requirement is that each complete accountable document be enclosed in front and back cover sheets, with the front sheet showing clearly all items listed in general principle 1 of Section III. The operator

can readily do this, deriving the other needed information besides classification from the audit log. In some instances, paragraph markings may also be appropriate, but, as has been mentioned earlier, such markings are deemed to be beyond the scope of the marking policy under consideration.

Non-paged output in this group will typically consist of printed interactive dialog, or possibly X-Y plots performed on continuous roll paper. The cover sheet requirements for this sub-class are identical to those for paged output. However, the classification markings internal to the material must be different, since no clear notion of "page tops and bottoms" exists at the time when the output is being produced. To avoid the loss of the advantages inherent in the pre-printed forms idea, a different type of pre-marking offers the best solution. In particular, half-tone pre-printed classification markings, appearing once for each imaginary "half-page" (i.e., about once every 6 inches or 15 centimeters), would provide an adequate interim indication of device and material level without unduly interfering with the legibility of randomly placed printing. When the material is actually removed from the device, it can be either divided into pages and marked properly or designated for destruction as classified waste.

Group 1B

Non-removable material that is directly human-legible typically consists of output from displays and from indicators. The actual distinction between displays and indicators will prove important to the consideration of multi-level policy, but for now it is irrelevant. Non-recording audio output would also fall in this group of I/O techniques.

The unique attribute of classified information revealed by this group of output mechanisms is that it need only be labeled with a single indication of the current classification level and category set of the output device, provided that the indication can always be readily viewed by the operator. No further markings are necessary because the output is not a document; it cannot be handed over to uncleared individuals, and it cannot even be circulated among cleared personnel. Only the sensitivity level need be indicated, so that the user may know with whom the displayed material may be discussed. Of course, any classified documents generated through the use of information obtained from this group of devices must be treated in the normal fashion; the audit log entry for the file that has been accessed will contain the extra details needed to properly mark the generated document.

What is needed, then, is a sign attached to the I/O device, indicating the current sensitivity level of operation. Ideally, the sign should not only bear a printed legend, but it should be color-coded as well. In passing, it may be noted that the value of color-coding would be greatly enhanced by the adoption of a global uniform color code for the various classification levels. Also, the means by which the sign is attached should be substantial enough that detachment requires some conscious activity, such as removing a few screws or undoing some frame latches. In this way, casual sign replacement can be avoided, and the probability of erroneous labeling can be reduced.

Group 2A

Some forms of removable output material are intended for human use without further digital processing, yet they cannot be used except in conjunction with some device. Two primary examples of this group of output items are microfilm on reels and audio tape recordings. For such material, a security marking policy must address two types of requirements, those for the physical material itself and those for the contained information that will eventually become directly human-legible.

Specifying the latter type of requirements is a task which brings to light the most fundamental difference between unilevel and multilevel policy. It would clearly be absurd to expect the output device operator to do any sort of marking on the contained information of material in this group. Therefore, the only reasonable approach is to insist that any such markings be applied by the computer system itself. Required computer-generated security labels must be accurate and trustworthy, since the actual physical protection (or lack of it) that is provided for information may depend upon them. The various issues involved with the generation of trustworthy labels will be dealt with in detail later in this paper. At present, it will suffice to say that producing these labels is a non-trivial matter involving some significant costs. For this reason, fundamental policy will dictate that computer-generated security labels will only be mandated when the circumstances of multilevel I/O make them indispensable; the labeling of contained information will not be required for cases of unilevel I/O. This does not prohibit uncertified system software or applications programs from applying security-related indications, even to unilevel I/O material. Any such actions, however, will neither satisfy output marking requirements nor affect input verification procedures. Furthermore, substantial management diligence will usually be needed to avoid the gradual development of a reliance upon such untrustworthy labels if they are allowed to exist.

The physical marking requirements are directed primarily at the reels, cartridges, envelopes, etc., that are normally used to encase Group 2A items. The general principle 1 markings are required only on these packages, since it tends to be impractical to place them directly on the output material. Gummed labels are typically used to apply the markings, with blank spaces provided on the labels for those details which vary from item to item. To ensure that the correct labels are applied, it is necessary to require that signs be placed on devices in the manner described for Group 1B output. Finally, it is generally advisable to color-code the item packages, so that a correspondence can be maintained between packages colors and device sign colors.

Group 2B

This group merely represents a meaningless intersection of the two group dimensions. If output material is not removable from its output device and yet not due for any further digital processing, then any legibility-inducing equipment required must be incorporated into the output device itself, placing the output technique in Group 1B. Otherwise, the output method would be useless. Clearly, then, no marking policy need be specified here.

Group 3A

The material encountered in this class is removable and normally read only by digital data processing equipment. It includes two basic types of media, perforated and magnetic. Examples of perforated media include fan-folded paper tape and punched cards, while the familiar types of magnetic material are tapes, disk packs, and floppy disks. For all of these items, computer-generated markings for the contained information constitute a matter of concern. However, the fundamental policy of not requiring such markings for unilevel output products, discussed in an earlier subsection, still applies here.

In the area of physical, human-legible markings, further discussions are still appropriate. The magnetic media tend to be very similar in nature to Group 2A material, and so the same marking policy, which requires labels for item containers and signs for the devices (and which advises the color-coding of containers), can apply. Perforated media, on the other hand, tend to lend themselves much more readily to direct physical marking. In fact, two distinct surfaces are generally available for applying markings: an initial, non-perforated area of item surface (e.g., a length of leader tape or a header card) and a composite surface made up of closely packed medium edges. Contemporary practice, as discussed earlier, seems to suggest that all markings required by general principle 1 of

Section III be applied to the initial surface, while the classification and category set of the item are stamped on the edge surface that is most commonly viewed. This seems to be a very reasonable policy. However, it must be augmented to the extent that signs indicating sensitivity level be required for the output devices. Also, special provisions, calling for the application of a full complement of markings to any segment of an accountable item which is separated from the whole, must be retained. Finally, the color-coding of actual perforated media seems desirable, though it may prove impractical in certain instances.

Group 3B

Unlike Group 2B, the intersection represented by this group is not meaningless. Though any inseparable data represented within one computer system would either be strictly internal or be directed towards an I/O device in some other class, a data message directed to a remote computer system via a network link would have to be thought of as being within Group 3B. However, any unilevel link may represent merely one step in a multilevel path of links, or it may be that a unilevel-at-both-ends path actually uses a few multilevel links. Therefore, any discussion of the requirements for this case will be deferred to the section on multilevel marking policy.

INPUT VERIFICATION POLICY

Groups 1A, 2A, 3A

The obvious application of the two marking policy objectives is to the placement of proper markings on generated output. In a data processing installation, the production of output is a highly visible activity, and it is very easy to imagine that the security marking for this heterogeneous mass of material must be subject to some sensible discipline. Another major installation activity, involving as much or more data but yet much less visible, is the creation of on-line data files. Modern virtual memory systems permit this activity to operate on an extremely large scale. The on-line files created in this way are directly analogous to the output documents produced by the more obvious reverse process. The same marking policy objectives apply to these files as apply to output documents, and it is input verification policy that relates these objectives to the creation of on-line data files.

The legibility dimension is not relevant to removable input; it is conceivable that any form of removable information-laden material could be used for computer input, and the source medium is not important once the data is in permanent on-line storage. Therefore,

all forms of removable input material may be dealt with as a single group. Furthermore, it may be assumed that any item which can be used for computer input might have been produced as computer output. This makes the specification of required markings for input material quite easy; to be accepted as input, an item must be marked as if it had been generated as output. Also, each input device must carry a sign indicating current sensitivity level, just as most output devices must do.

The actual verification procedure comes in three stages. The first stage is operator inspection of the material. This inspection should ensure that all required physical markings are present, and that those markings do not indicate that the item in question is multilevel when the input device is operating in a unilevel mode. Those markings that indicate the multilevel nature of an item will be detailed in the next section.

The second stage of input verification involves checking the operating sensitivity level of the input device to be used. This checking is done by comparing the sign on the device to the markings borne by the item. Extensive use of color-coding can cause any errors committed at this point to become very noticeable.

Both of the first two stages must take place before any actual input operation can proceed. The third stage takes place after the input operations have been completed, but just before the operator returns the input item to physically secure storage. At this time, the operator must check all audit log entries that have been generated as a consequence of input from the item. This responsibility includes both verifying the item-related information that the system has already entered and providing the details needed to complete those entries which the system could not complete on its own.

In particular, though there is only one possible file classification in a unilevel system, such information as downgrading category, declassification date, and source of classification authority must be entered. These details, among others, will have been included in the external physical markings of the input item. Upon completion of this third stage, input verification has been accomplished and the item may be returned to secure storage.

Group 1B

For input, this group includes all of the myriad forms of manual data entry, as well as voice data entry. Since the input data has no physical substance, there is no way it can actually

be marked before it is presented to the system. Therefore, the thrust of input verification policy for this group is simply to ensure that the user entering data knows at what sensitivity level the input device is operating. This can readily be done by taking advantage of the fact that, just because the data entered has no real substance, systems generally echo such input back to the user. The echoed data is at the same sensitivity level as the input, and furthermore it is governed by the output marking policies of Group 1A or 1B, depending on the nature of the device. The marking of the echoed information suffices to inform the user of the level to which the system will protect the input data. Of course, when a user's activity results in the creation of a new on-line file, he or she must complete the audit log entry by entering such details as downgrading category, source of classification authority, etc.

SECTION VII

A GENERAL SECURITY MARKING POLICY IN THE PRESENCE OF MULTILEVEL I/O

INTRODUCTION

When a particular input/output device is running in a unilevel mode, the appropriate security marking policy is neither difficult nor expensive to carry out. For this reason, unilevel I/O is an attractive operating option. Furthermore, as was demonstrated in Section V, it is often the only operating mode which may be permitted. It does, however, involve one serious penalty: if material of different sensitivities must be processed under time constraints which do not permit sanitization or reconfiguration for each level change, then multiple identical I/O devices, one for each level, must be used. Situations can easily arise in which multiple levels must be handled, but the use of multiple devices represents a prohibitive expense and logical security enforcement does not require unilevel device operation. Under these circumstances, it is necessary to turn to the alternative operating option of multilevel I/O.

It has previously been pointed out that there are three situations in which logical security enforcement will permit multilevel input/output. Any pure output device, i.e., any output device with no functioning input capability, may handle multilevel traffic. Devices that do perform input functions may process multiple levels without human intervention, provided that those devices are controlled entirely and exclusively by trustworthy software, and that the input material itself includes trustworthy, machine-legible labels indicating the data sensitivity level. Finally, any terminals used for interactive processing may be operated in a special multilevel mode which calls for trustworthy software to set a device's operating level, uncertified software to control the device's operation at the level, and then control to be returned to trustworthy software whenever level changes are to be effected. The marking policies for specific devices will make reference to the particular situation which can justify multilevel operation for the device in question.

Two points of general underlying policy can be established before the various I/O technique groups are explored. First of all, every device operating in a multilevel mode must display two attached signs, one giving the device clearance (exactly like the level-indicating signs of unilevel policy) and one indicating the fact that the device is engaged in multilevel operation. Secondly,

the differences in level between successively produced accountable documents, and, when appropriate, the differences between distinct sections of individual documents, must be indicated by trustworthy labels generated in the course of output production. The techniques for ensuring the trustworthiness of computer-generated labels will be discussed in the following section. In the remainder of this section, multilevel output marking policy and input verification policy for the various previously defined groups of I/O techniques will be reviewed so that the differences from unilevel policy can be examined in detail.

OUTPUT MARKING POLICY

Group 1A

Multilevel operation has no impact on the cover sheet requirements for human-legible, removable output material. However, the use of pre-printed forms is clearly no longer appropriate, since the level of the material to be printed can no longer be predicted with accuracy. For line printers using continuous, fan-folded forms, the pre-applied classification markings at the top and bottom of every page must be replaced by computer-generated sensitivity level labels printed in the same places. In addition, the system-generated pages which separate the individual print-outs should carry large banner markings giving classification and category set, so that the separation of material at different levels may be facilitated. These banner pages cannot take the place of cover sheets, since frequently several print-outs at different levels will be combined to form a single accountable document, and the cover sheets must show the maximum overall sensitivity level of such an enclosed item.

Line printers obviously qualify for multilevel operation as pure output devices. Those terminals which produce non-paged printed interactive dialog, on the other hand, must have their control switched between trustworthy and uncertified software in order to function in this more sophisticated fashion. Non-paged material, then, does not need constant security labeling, but rather only indications of level setting, printed when that setting takes place. Since the indications are printed as hard copy, they can be referred to whenever a user has doubts about the current operation level, thereby satisfying the requirement that they always be in view. However, it must also be recognized that interactive dialog consists of two components, echoed keyboard input and system-generated output. The sensitivity level of the latter may change without any change in the input level if, for example, the user requests a listing through the terminal of a file at a sensitivity level lower than the current

operating level of the terminal. (Analogous changes in the sensitivity level of the input are not possible because logical security enforcement precludes changes in the input level that are not accompanied by changes in the actual interactive processing level, i.e., changes in the terminal operating level).

These considerations suggest a policy in which the computer system is made responsible for printing two types of security messages, those for the keyboard and those for the printer. While keyboard messages are properly dealt with under input verification, it may now be stated that they will give the current operating sensitivity level just after log-on, just before log-off, and both before and after every operating level change which occurs during the course of an interactive session. These messages will in general suffice for output marking as well. However, each time a change occurs in the sensitivity level of the output without a corresponding change in the interactive processing level, then the system will be responsible for printing out special messages indicating the initiation and termination of the output level change. Finally, the interactive user is expected to divide into pages and properly mark all material not designated for destruction as classified waste, just as in unilevel policy.

The hard copy X-Y plotter, whether fed with single sheets or continuous rolls of paper, presents a different sort of problem. Here again, multilevel policy would dictate the replacement of pre-printed classification markings with trustworthy computer-generated labels. However, in Section VIII, which will deal with the implementation of trustworthy labels, it will be shown that applying such a policy to this particular case may prove so difficult as to be impractical. If this is true, then X-Y plotters, though undeniably qualified as pure output devices, must nevertheless be considered unsuitable for multilevel operation.

Group 1B

The distinction between displays and indicators, unimportant for unilevel policy, now assumes a pivotal role. A display is an area, controlled as a unit, that can show several letters, numerals, lines, or condition indications at once. Its output is characterized by the ability to change location within the area, and to do so without any changes in meaning. An indicator, controlled as a unit, can show only a single letter, numeral, or condition indication. A cluster of indicators may combine to show composite information, but they do not constitute a display unless they are controlled as an integrated unit and possess the information movement property. Thus, a CRT (cathode-ray tube) screen is a display, but a row of lights showing a memory address in binary notation constitutes an indicator cluster.

The policy impact of the distinction is that displays can reasonably be expected to show security information continuously along with output data, while indicators and indicator clusters cannot be expected to do this. For multilevel operation, then, the sign that used to show the device's current classification and category set, formerly attached to a display device, is now replaced by an analogous indication exhibited continuously by the display itself. Maintaining the indication continuously is necessary to satisfy the always-in-view requirement. If a device includes an indicator dedicated to showing sensitivity level information, then that indicator can serve as the replacement for the sign that used to define the current level. (It should be kept in mind that signs are still required for showing the current device clearance and the fact of multilevel operation.) However, if a device includes only indicators, none of which are dedicated to showing sensitivity level information, then that device can only be allowed to function in a unilevel mode.

All of the above applies primarily to devices that qualify for multilevel consideration by dealing exclusively with data output. However, many displays and indicators appear on devices that include some means for data input and that must qualify for multilevel consideration by having control switched between trustworthy and uncertified software. As happened in Group 1A, the real output may become combined with echoed input, and the problem of output changing level without the input doing the same arises again. When this occurs, two distinct display indications, one for input and one for output, or two separate dedicated sensitivity indicators, or a combination of a display and a sensitivity indicator, must be used. A device capable of input and output, but possessing no display and only one dedicated sensitivity indicator, is inadequate for multilevel operation.

Non-recording audio output also belongs to this group of I/O techniques, though outside of the area delineated by the display-indicator distinction. As was the case for the X-Y plotter, implementation considerations will be shown to argue against audio security indications, so that multilevel operation appears to be precluded for this technique, also. While the use of a dedicated visual security indicator might seem to represent a potential solution, the prospects appear dubious, primarily because a user is not likely to continually pay strict attention to a visual indicator while trying to concentrate on the audio output itself.

Group 2A

Those requirements for the marking of encasing packages of Group 2A material that were established under unilevel policy still

apply in the multilevel case. Two new requirements must be added, though. One involves placing an additional label on each package which states that the enclosed item is multilevel and contains internal labels. The other requirement is that those internal labels mentioned by the additional sticker must actually be present.

The nature of the required labels for contained information is actually quite easy to specify. Any human-legible data that results from the action of an appropriate device on the material of this group must be directly analogous in form to some type of material included in Group 1A or 1B. The most reasonable policy, then, is to extend the provisions for multilevel markings, already established for the two directly human-legible groups, to cover the analogous information contained in Group 2A. In the case of a microfilmed print-out, for example, each print-out page must be marked at the top and bottom with a sensitivity level indication generated by the system, just as if that page were a full-sized one produced by a line printer. However, the analog of an audio tape recording's contained information is direct audio output, and no multilevel requirements were established for that output technique due to implementation considerations. Thus, a policy extension leads to the conclusion that computer-generated audio tape recordings also do not represent a medium suitable for multilevel use.

A note of caution must be added concerning the use of color-coding for multilevel Group 2A material packages. The color of the sign attached to an output device now represents that device's operating clearance, not necessarily its current operating level. Items produced by a multilevel device running at a given clearance may not contain any information at the clearance level, and therefore may possess an overall lower classification. For this reason, it becomes necessary for operators to base all of their markings, including the overall item-classification marking, exclusively upon the audit log information. A microfilmer running multilevel and Top Secret-cleared, for example, might well produce a reel of microfilm carrying information classified no higher than Confidential. The appropriate can for this reel would be color-coded for Confidential, and would not match the Top Secret color-coded clearance sign attached to the filmer. Obviously, then, this situation must be clearly understood by operations personnel, or else color-coding should be abandoned altogether for this sort of multilevel operation.

Group 2B has previously been shown to be a meaningless intersection of the two group dimensions.

Group 3A

Under unilevel policy, provisions for the physical marking of both perforated and magnetic material of this group were established. As was the case for Group 2A material, these provisions can be retained under multilevel policy. Of course, the same problem with color-coding exists, so that the audit log becomes the sole source of information on which to base markings. Furthermore, the use of color-coded perforated media for multilevel devices must always be prohibited, since the level of the data to be punched cannot be accurately predicted.

The physical markings required must be increased to include an additional gummed label, one that states that the item is multilevel and includes contained labels. This extra requirement was also mentioned in connection with Group 2A material, but it is much more important here. Multilevel material of Group 3A is more likely to be used for actual multilevel input than any other type of material. Such input is not to be permitted except when trustworthy software alone controls the input device, and it is only by means of the multilevel identification sticker on potential input material that device operators can enforce this policy.

When devices in this group can qualify for multilevel operation, either by being pure output devices (as will often be the case for perforated material production) or by being controlled exclusively by trustworthy software (as will generally be required for magnetic material production), and are in fact operated in that mode, then their output must include trustworthy, machine-generated sensitivity level labels. When an item accountable to a single user is produced, information at different levels must be written in distinct records or analogous units, with a label included before and after each unit so that the units are delimited. After the appropriate physical markings have been applied, the item may be delivered to the user.

Some tapes and disks will contain material stored as part of the on-line file storage system, and such material will typically be accountable to a number of users, in addition to being multilevel. While the internal labeling requirements for these special items are not different, the physical markings, both on the items themselves and on the I/O devices that produce them, must be unique and distinct enough to ensure that these items are never circulated outside of the machine room where they were generated. Without such special provisions for I/O material bearing on-line files, individual user accountability cannot be maintained.

Group 3B

This group includes data messages sent from one computer system to another via a network link. It will be assumed that the link itself possesses a clearance (probably related to its degree of cryptographic protection), and that it handles messages at a variety of sensitivity levels at and below its clearance level. The link qualifies for multilevel operation by being controlled exclusively by trusted software, though uncertified software may certainly prepare messages, request their transmission, and be given the text of incoming messages. Finally, it will be assumed that each individual message contains data at only a single level, and that transmissions of material at multiple levels can readily be broken down into discrete single-level messages.

With this background established, developing the marking requirement becomes relatively easy. A secure computer which sends a classified data message through a network link would be required to make an audit log entry for the event, just as for any creation of accountable output. Much the same applies to the machine which receives the message. Indeed, the security information sent with the message must be sufficient to allow the receiving system to properly complete its own audit log entry. The most simple and direct way to insure that enough information is sent would be to require that the sending system transmit, as a security label, the contents of its own audit log entry for the event of the message's transmission. The technique by which this information is incorporated into the transmitted data depends heavily upon the network communications protocol used, and so any further specific details lie beyond the scope of a marking policy.

INPUT VERIFICATION POLICY

Groups 1A, 2A, 3A

The role and the importance of input verification policy were discussed in the exposition of unilevel policy, and that discussion remains valid in the context of multilevel policy. The idea of the legibility dimension's irrelevance to removable input is still applicable, and the policy of requiring that input items be marked as if they had been output items can be carried over without change. Some modifications must be made, however, to the specifics of the three stages of actual verification procedure. The key to those changes is that multilevel input of material in this composite group can be permitted, and can only be permitted, when the input device is controlled exclusively and entirely by trustworthy software.

Under unilevel policy, the first stage, which involves operator inspection of the input material, called for the rejection of multilevel material. Clearly, when multilevel input is expected because trustworthy software controls the device, such rejection is no longer appropriate.

The second stage used to require that the operator check the level of the input material and make sure it corresponds exactly to the operating level of the input device. This was done by comparing the material markings with a sign attached to the device. Now, when the device bears a sign saying that it is operating in a multilevel mode, the level indicator sign merely represents the device clearance. The level of the material need not correspond precisely to the sign-indicated level; it may be lower. However, the operator must still make sure that no material fed to an input device exceeds that device's clearance level.

The requirements of the third stage, in which the operator checks and completes audit log entries, remain unchanged for instances of user-requested multilevel input. However, multilevel input of material in this composite class may consist of actions taken on behalf of the on-line file storage system. Material involved in file storage transactions can only carry data that had already been accounted for on the "creation" audit log, and therefore input of this special type will not result in any new audit log entries.

Group 1B

Multilevel manual data entry can generally be expected to take place only on I/O devices set at a level by trustworthy software and then controlled at that level by uncertified software. As was the case for unilevel policy, input verification for this sort of material can consist entirely of appropriate output labeling of echoed input. The multilevel output marking policies of Groups 1A and 1B both make adequate provision for the indication of echoed input sensitivity level. The requirement for the user to perform audit log entry completion when appropriate remains unchanged.

Group 3B

The markings required for acceptance of a classified data message received through a computer network link are the same as those required for sending it. Full responsibility for entirely completing the audit log entry for each incoming message and then placing the message in appropriately protected storage rests, of course, with the receiving computer system.

SECTION VIII

THE IMPLEMENTATION OF TRUSTWORTHY LABELS

INTRODUCTION

The issues that have been defined and discussed in the preceding sections have been policy issues, specifically, the development of security marking application and verification requirements that satisfy the two original policy objectives. These are questions of what is to be done in a variety of situations. Questions of how established policy is to be carried out fall into the category of implementation issues. The nature of many issues of this kind depends upon the properties of a particular system design, so an exhaustive review of implementation questions is not possible in the absence of a specific design. However, there are certain problems pertaining to the trustworthiness of machine-produced security labels that can be discussed in a general context. This section will be devoted to dealing with them. Of course, it must be recalled that concept of computer-generated security labels is relevant only to multilevel marking policy.

CORRECT ISSUANCE OF LABELS

The first requirement for trustworthy labels is that they be issued correctly by the system. Correct issuance has two aspects; the labels must consist of correct information, and they must be issued at the proper time. This need for correct function performance puts label issuance in the same class as other security enforcement functions. That is, label issuance must be performed by certified software as part of an integrated security enforcement mechanism.

The incorporation of input/output operations into a security enforcement mechanism is discussed by Burke. [4] According to his concept of operations for handling I/O in a secure computer, the I/O controls must perform three major operations for every input/output transfer or sequence of transfers. These operations are authentication, controlled attachment, and controlled operation.

The objective of authentication is to establish the identity of the user or medium at the terminal or device drive, respectively. This process permits the appropriate sensitivity attributes to be associated with the data source or data sink in question. Controlled attachment usually refers to the logical or software attachment, i.e.,

the "making known" to the computer system of an external candidate for I/O transfers. In some cases, however, this process may include the verification of the hardware link as well.

It is the process of controlled operation that must include provision for label issuance. The primary objective of this aspect of I/O operations is to insure that the actual execution of a transfer does not cause a change in the device attachment. For I/O devices operating in a unilevel mode within a multilevel system, the controlled operation function need not have any other concerns because no labels are produced along with the output. For multilevel devices, on the other hand, the correct issuance of labels becomes a major responsibility of the controlled operation function.

The functions that control the operation of multilevel devices will clearly be more complex, and therefore will be more difficult and expensive to program and certify than unilevel control functions. This problem is compounded by the fact that these multilevel routines must be especially tailored to the particular characteristics of each individual input/output device. Such special I/O control routines represent the first cost of generating trustworthy labels. The second cost is incurred in reliably conveying the correctly issued labels to the output material itself. This task can be accomplished in one of two basic fashions.

One technique involves the use of a reserved I/O channel, dedicated exclusively to carrying security labels. Such a dedicated channel might feed security information to a dedicated sensitivity level indicator, for example, or possibly to a special separate printing element added to a line printer and dedicated to the production of distinct security labels. This technique cannot be applied to all output methods (e.g., multilevel microfilmed output), and it often cannot even be used with suitable methods except at prohibitive cost. The latter is especially true if the dedicated I/O channel becomes in reality an encrypted, lengthy, terribly under-used communications line. When a separate channel for security labels cannot be used, for whatever reason, then the only available option for conveying the labels to the output material is to mix them in with the actual output data and send them over the same I/O channel. Under these circumstances, label spoofing and label alteration both become threats that must be neutralized if the labels are to remain trustworthy.

PRECLUDING LABEL SPOOFING

Spoofing is a term used to describe the unauthorized generation of imitation security labels by uncertified software. These imitation labels can be included in the output data with which the

genuine labels are mixed. Through this technique, it is possible to misrepresent the sensitivity level of an output document, with the result that cleared personnel may unwittingly compromise classified material by handling it in accordance with the false labels.

It is important to distinguish between the general problem of label spoofing and the more specific problem of interactive system spoofing. The scenario of the latter involves a malicious user who writes a program that exactly imitates the responses of a particular interactive system. Another user may then think that he is dealing with a legitimate system when in fact he is only feeding data to the malicious user's program. There are several simple ways for a user to ensure that he is actually dealing with a legitimate interactive system, e.g., he can hang up and call back, or he can use a special key to generate a hardware interrupt. However, these techniques do not adequately deal with the general label spoofing threat, since the problem typically applies to such non-interactive devices as line printers and tape drives.

There is only one way to preclude general label spoofing, that being to reserve some capability of the I/O device in question for use by certified software only. This capability must then be included in every security label in such a way that it delimits the label, and the I/O control routine must censor the capability out of any output material presented by uncertified software. For example, in the case of a line printer, a particular special character, perhaps the number sign (#), could be reserved. The large labels required on the pages which separate print-outs could be made up entirely of reserved characters, and the page labels could be surrounded by borders of them. It will be recalled that borders of asterisks are now used to differentiate between security labels and text on teletype messages (see Section III, Figure 3).

Of course, the value of the reserved character is that its presence is checked for in messages prepared by uncertified software and censored out of those messages. Any attempt to spoof security labels will clearly appear as nothing more than an attempted spoofing. However, a price is paid for this achievement. In addition to losing a character from regular use, the speed of the I/O device may be reduced due to the checking of untrusted messages which must be carried out. For the line printer, the check is a simple one, merely a comparison of each character with the reserved character to make sure that they are not the same. This could be performed very rapidly, with relatively little impact on printing speed. However, the same does not hold true for all I/O techniques.

Two examples of much less simple cases were referred to in the discussion of multilevel output marking policy. One is the X-Y plotter, which draws lines rather than printing discrete characters. The other is audio output, which consists of a succession of tones. In both of these cases, there is no particular capability that could be easily recognized by device operators and yet easily and quickly checked for in untrusted messages. Thus, spoofing could not be precluded without a serious impact upon device performance, not to mention the development of a complex, difficult-to-certify checking routine. These devices, then, and others with the same problem, cannot be used for multilevel output, since it is impractical to enable them to produce trustworthy security labels.

PRECLUDING LABEL ALTERATION

There are some cases of I/O in which even correct and non-spoofable security labels cannot necessarily be trusted. This condition results from the fact that the technique involved is one which permits easy and undetectable erasure of data. Two examples of such erasable media are magnetic tapes and CRT screens. When easy erasure is possible, uncertified software can, by adding data to the output stream, gain access to the area within the reserved character border, erase the original contents, and enter new, false labels. This technique is called label alteration, and it is quite distinct from label spoofing, which involves the generation of entire false labels where none had previously existed.

As with spoofing, there is only one basic way to preclude this compromise risk. In this case, it involves preventing uncertified software from accessing any erasable security label area. For example, in the case of multilevel magnetic media, the trusted control software can readily ensure that each record begins with a label, is written in one direction only, and ends with a label.

The CRT display represents a more difficult case in terms of controlling label alteration. Several CRT's come equipped with a "format" mode, in which some fields on the screen are fixed and only the variable fields may be altered. This feature might reduce the task of certified software to merely ensuring that uncertified software cannot kick the display out of the special mode. Other devices of this type maintain registers indicating the coordinates of the cursor position. Monitoring those registers could be a reasonable task for a certified routine. However, some CRT's do not provide any satisfactory technique for controlling cursor positioning (that is, in the sense of preventing the cursor from reaching a certain position), and these must be deemed unsuitable for multilevel use with a secure system.

Once the three implementation considerations that have been mentioned in this section are taken into account, trustworthy labeling becomes feasible. The operator who must use the labels for guidance can be sure that they were produced correctly in the first place, that they are not fake labels generated by uncertified software, and that they have not been altered after being produced. No further doubts about the validity of such labels can be entertained, and the existence of trustworthy labels allows important aspects of the general security marking policy to be implemented.

SECTION IX

SUMMARY

General security marking policy has been introduced by describing its two objectives. One objective was to create a consistent and generalized basis for establishing the security marking requirements for a wide variety of input/output techniques. The other was to ensure that the security attributes of classified material would be conveyed accurately across the boundary between physical and logical security enforcement. A background of current regulations and existing policies was described, and then a general policy was presented in detail.

Two key concepts, one directed at each policy objective, led to the evolution of a detailed policy. The basis for marking requirements was developed by defining a set of groups of input/output techniques based upon the two group dimensions of removability and legibility. Ensuring the accurate transference of security attributes was greatly facilitated by drawing a distinction between unilevel I/O, in which operator intervention accompanies each level change and no machine-generated labels are required, and multilevel I/O, in which changes occur without intervention and labeling by the computer become necessary.

Also involved in the development of a general marking policy were two significant secondary concepts. It was shown that the functions performed by logical security enforcement mechanisms permit multilevel output but preclude multilevel input except under certain very specific circumstances. With respect to the implementation of machine-generated security labels, it was explained that if these labels are to be trustworthy, they must be generated by certified software. Furthermore, they must be conveyed to the actual output either on a separate, dedicated channel or in such a fashion as to neutralize the risks of label spoofing and label alteration.

However, the most important single aspect of the policy that has been described is that, while fully achieving both objectives, it still leaves the most crucial decisions in the hands of each installation manager. Specifically, the question of whether to operate I/O in a unilevel or multilevel mode can be decided independently for each individual input/output device, based exclusively upon local perceptions of the cost-vs.-throughput tradeoffs involved. The general policy merely endeavors to describe the choices; the actual choosing is done by the facility management. In this way, essential flexibility is combined with the degree of rigidity necessary to achieve the objectives of general security marking policy.

REFERENCES

1. Anderson, James P., Computer Security Technology Planning Study, ESD-TR-73-51, October 1972.
2. Bell, D. Elliott and LaPadula, L.J., Secure Computer Systems: A Mathematical Model, ESD-TR-73-278, Vol. II, November 1973.
3. Bell, D. Elliott, Secure Computer Systems: A Refinement of the Mathematical Model, ESD-TR-73-278, Vol. III, April 1974.
4. Burke, Edmund L., Concept of Operation for Handling I/O in a Secure Computer at the Air Force Data Services Center (AFDSC), ESD-TR-74-113, April 1974.
5. Department of the Air Force, Air Force Message Management Program, AF Manual 10-2, July 1970.
6. Department of the Air Force, Management of Data Processing Equipment, AF Manual 171-9, July 1970.
7. Department of the Air Force, Information Security Program, AF Regulation 205-1, February 1973.
8. Department of Defense, Regulation Governing the Classification, Downgrading, Declassification, and Safeguarding of Classified Information and Material (Short Title: DoD Information Security Program Regulation), DoD 5200.1-R, July 1972.
9. Department of Defense, Techniques and Procedures for Implementing, Deactivating, Testing, and Evaluating - Secure Resource-Sharing ADP Systems, DoD 5200.28-M, January 1973.
10. Electronic Systems Division (AFSC), Deputy for Command and Management Systems (MCI), ESD 1974 Computer Security Development Summary, MCI-75-1, December 1974.
11. Whitmore, J. et al., "Design for Multics Security Enhancements," Honeywell Information Systems, Inc., ESD-TR-74-176, December 1974.
- k2, Honeywell Information Systems, Inc., WWMCCS Security User's Guide, Series 6000 Software, WWMCCS 1, Rev. 1, June 1972.

13. Joint Technical Support Activity, "Security Features of the WWMCCS GCOS-III System," Operating System Technical Bulletin Number 72-OS-03, May 1973.
14. Weissman, C., "Security Controls in the ADEPT-50 Time-Sharing System," AFIPS Proceedings, Vol. 35, FJCC, 1969.