NISTIR 4362

SECURITY LABELS FOR OPEN SYSTEMS AN INVITATIONAL WORKSHOP

Noel Nazario
Chairman

U.S. DEPARTMENT OF COMMERCE
National Institute of Standards and Technology
National Computer Systems Laboratory
Gaithersburg, MD 20899

U.S. DEPARTMENT OF COMMERCE
Robert A. Mosbacher, Secretary
NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY
John W. Lyons, Director

NIST

NISTIR 4362

SECURITY LABELS FOR OPEN SYSTEMS AN INVITATIONAL WORKSHOP

Noel Nazario
Chairman

U.S. DEPARTMENT OF COMMERCE
National Institute of Standards and Technology
National Computer Systems Laboratory
Gaithersburg, MD 20899

June 1990

U.S. DEPARTMENT OF COMMERCE
Robert A. Mosbacher, Secretary
NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY
John W. Lyons, Director

Table of Contents

Abstract

On May 30 and 31, 1990 the Protocol Security Group at NIST held a Workshop on Security Labels. Thirty-Five representatives from the U.S. Government, industry, and the United Kingdom gathered for two days to discuss security Labels for open systems. The discussion went from the generalities of labels in "end systems" to the more specific issue of labels in secure Open Systems Interconnection (OSI). The information shared during the two days of discussion is documented in these proceedings.

Key Words: Government open systems interconnection profiles; Integrity; Network; open systems interconnection; security labels; trust

Papers are contributions of the authors and do not necessarily represent the views of NIST.

Workshop Report

Robert Rosenthal (NIST) welcomed the attendees. He indicated that the group would first look at labels in general and then focus on the role of labels in Networks. Mr. Rosenthal expects to incorporate the output of the workshop into the U.S. Government Open Systems Interconnection Profile (GOSIP) FIPS PUB 146 (3). He then introduced the chairman of the workshop, **Noel Nazario**.

Mr. Nazario thanked everyone for their written contributions. After reviewing the workshop agenda and asking all attendees to introduce themselves, he introduced the technical presentations.

Dr. Dennis Branstad (NIST) gave a top down view of the labelling problem. He talked about the European Computer Manufacturing Association's (ECMA) "Security in Open Systems Framework, TR 46", and the ECMA Data Elements and Service Definitions". Dr. Branstad talked about security policy domains and users within the security domain. He described the relationship of labels between the domains and the requirement of domain rules. The existence of these domains introduces the problem of "across-domain" communication. This requires a method of translation. From his perspective, the purpose of the workshop was to look at the content of security labels and come up with a common generic format. This generic format should be tailored to different domains using registration authorities to maintain consistency.

Dr. Branstad also mentioned that ECMA has associated attributes with users, subjects, and objects by means of labels. The definition of the information in a label and of the defining authority are yet to be determined. NIST is in the process of studying the possibility of becoming a National Federal Registering Authority.

Three security services are associated with a labelling scheme: confidentiality, integrity, and availability (C, I, A). Some felt that since there are no rule-based mechanisms to enforce availability, it should be handled as a Quality of Service (QOS) attribute.

David Chizmadia (NCSC) talked about the purpose of NCSC's labelling guideline for end systems. In his presentation he talked about the needs of the user community. He also discussed the pros and cons of labels and their implementation.

Phil Mellinger (MITRE) gave a talk on the Blacker network front end (BFE) and how it makes access control decisions based on the Internet Protocol Security Option (IPSO) labels.

Some suggested the integration of a restricted routing service. Mr. Mellinger indicated that the BFE does not make routing decisions.

Other points made in Mr. Mellinger's presentation were:

- Policy should be established by the customer,
- Standardize while preserving flexibility,
- Determine label content,
- Use multiple authorities for a data label,
- Determine how end-systems should deal with authorities,
- Define a security policy and labels that fit RFC-1038, and
- Access rights should not be carried in the label.

John Linn (Digital) presented "Issues on the Commercial Internet Protocol Security Option (CIPSO)." He explained how the original IPSO is oriented towards classified requirements and does not satisfy the needs of the commercial user. Another point he made is the need to define labeling authorities to accommodate different security policies.

Ali Eshgh (SSDS) talked about a decentralized network security administration that requires an agreed upon set of standards to be useful and effective. A network passport concept was described where each domain is responsible for authorizing traffic. Each user within a domain has a token. Users also have a visa that is unforgettable.

Russell Housley (Xerox) raised the issue of avoiding availability, authorization and billing codes in security labels. He stated that availability is a Quality of Service. All this has to do with defining computer security but, should we shove it all under labeling?

Bill Maimone (ORACLE) talked about the database community's concern with end-system and network security labels as they apply in distributed processing. ORACLE's products use security label information directly from the operating system. Database labels need to be consistent and his company would like consistency in operating system labels as well. An application level standard should be encouraged. Mr. Maimone's position paper outlines the role of secrecy labels as used by a hierarchically subset database management system and the subsequent requirements for standardization of labels.

Dr. Robert Shirey (MITRE) reviewed the contents of his position paper that talks about various components of the Defense Message System (DMS) and the three implementation phases the DMS project.

John Woodward (MITRE) discussed information labels, a concept developed under the Defense Intelligence Agency/MITRE Compartmented Mode Workstation (CMW) project. By using information labels it is possible to track the classification history of data to provide safeguards and avoid overclassification. As part of his presentation Mr. Woodward showed an example of a language for defining label encodings that could be used in policy generation.

Gene Troy (NIST) talked about C, I, A as pertinent to security labels. Unclassified confidentiality is similar to classified problems. The question raised is, how to describe confidentiality?

Dennis Steinauer (NIST) reviewed security labels as related to the Portable Operating Systems Interface (POSIX). He stressed the need for a standard labeling mechanism.

Nick Pope (Logica, U.K.) talked about security label work in Europe. Some of the points he covered are:

- Security policy ID
- Classification
- Primary Mark
- Security Category (object identifier)
- Label information be carried in 2 ways, in a label or as a Quality of Service parameter.
- Security framework

Warren Schmitt (Sears Technology Services) presented the security needs of commercial institutions. Mr. Schmitt pointed out risk or exposure areas that businesses need to protect against the most. He also stated that the commercial community's concern about confidentiality, integrity, and availability is evenly spread and not necessarily focused on one or two of these.

Day one ended after Mr. Schmitt's presentation.

Dr. Stuart Katzke (NIST) summarized the relevant points of the previous day's discussion. He talked about the wide acceptance of a relationship between security labels and confidentiality but pointed out that their relationship with integrity and availability is still being debated. In addition he made the following points:

- Not all data units need a label,

- There is a relationship between the form of the label and the security domain in which the host resides,

- The function of the different OSI layers should be considered when defining a labeling scheme, and

- Trust is required on the label itself.

Russell Housley (Xerox) presented his views on security labels and their placement within the OSI architecture. He provided definitions for data security and security labels. Mr. Housley also described types of protection using fundamental security models. Integrity and confidentiality in labels were discussed along with reasons for treating availability as a QOS parameter. He ended with an analysis of the pros and cons of security labels at each layer.

Douglas Brown (Department of Energy) presented DoE's approach to IP Security Labeling which has been revised as a proposal for GOSIP (CLNP) labels. Mr. Brown also provided a review of the work done by the Trusted Systems Interoperability Group (TSIG) on the Commercial Internet Protocol Security Option (CIPSO).

Mr. Brown provided the following background information on the DoE effort:

> DoE chartered a Security Labeling Standard Working Group which adopted and extended the Revised IP Security Option defined by RFC 1038 in the following ways:
>
> - Used Basic Security Option without change (4 labels U, C, S, TS)
> - Adopted an additional Protection Authority Code (4).
> - Added a DoE protection Flag.
>
> A reason to choose this approach was that the DNSIX interface specification developed by the Mitre Corporation specified the use of the Basic Security Option to communicate security levels.

Mr. Brown also gave a description of the DoE basic security model and the justification for its use.

John Woodward provided background information on the Extended Security Option (ESO) and compartments in IP/CLNP labels. He explained the meaning of several acronyms and the processing of security labels in Intelligence Systems.

N. Vasudevan from IBM talked about a labeling approach for distributed systems. He covered end system labels moving all the way to security labels in open heterogeneous networks.

Open Discussion:

Noel Nazario opened a discussion to obtain a few points of agreement to be stated as output of the workshop.

- The overall scheme for security labels should identify country versions for security labels.

  Given that a unified labelling scheme for secure OSI would be presented to the international community by U.S. delegates, a provision has to be made for distinguishing between label versions for different countries. Such a field would be hierarchically expanded to identify registration authorities.

- Options 130 and 133 (Basic Security and Extended Security Options) should be enhanced with the TSIG's Commercial IPSO options.

- SP4 and CLNP should use the same kind of security label.

- NIST should be the Registration Authority for security labels.

- This group should review sections on security labels added to GOSIP by NIST.

Workshop Contributions

"Security Policies, Domains, and Labels", Dennis Branstad (National Institue of Standards and Technology)

Security Policies

Domains and Labels

Dr. Dennis Branstad

NIST Computer Science Fellow

Goals

- Support Specifications and Registration of:
    * Security Policies (P)
    * Security Domains (D)
    * Security Labels (L)
- Develop Standards for Distributed Systems
    * Trusted End Systems Supporting PDLs
    * Secure Communications Enforcing PDLs

Security Policies

- Specified by an Organizational Entity
- Responsibility of a Security Administrator
- Define a Security Domain
  * Example: DoD Classified Information Security Policy
  * Example: IBM Information Security Policy

Security Domains

- May have peer-to-peer relationships
- May have sub-domains
- Interdomain rules are required
- Mobile User Support
    * Originating (Source) Domain
    * Authenticating (Home) Domain
    * Destination Domain
- Each Domain has Security Labels

Non-Intersecting Domains

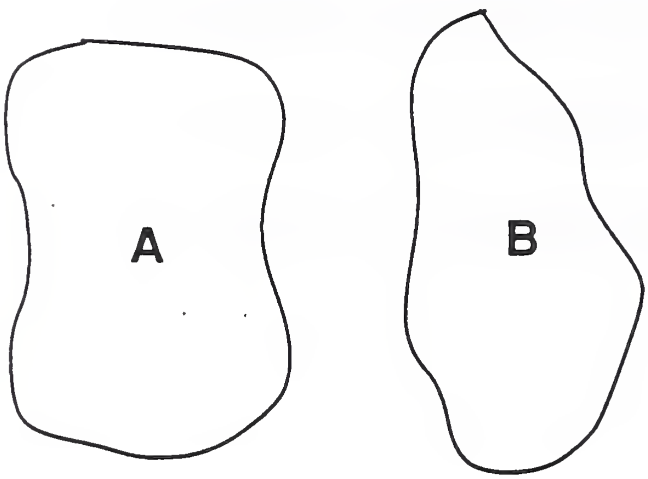

DoD Classified Financial Medical

Intersecting Domains

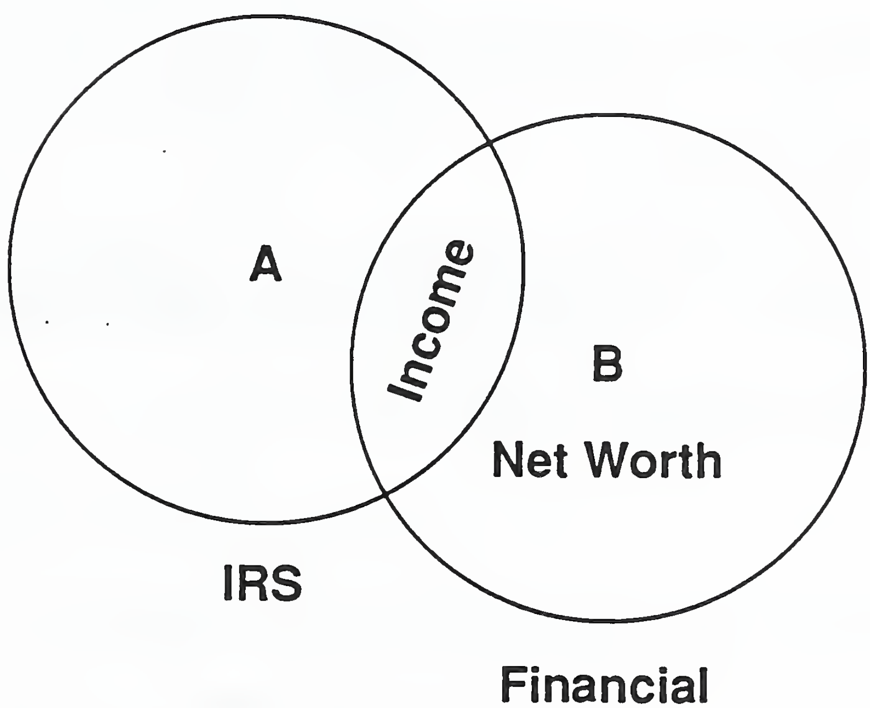

Security Labels

- Required to Enforce Policy in a Domain
- Must be Bound Securely to the Object (Data)
- Types of Labels
  * Subject Identity
  * Object Identity
  * Object Confidentiality, Integrity, and Availability
  * Subject Access Privilege Attributes (Clearances, etc.)

Typical Security Labels

- Typical Contents
    * Personal Identifier
        - Registration Authority
        - Name
        - Place of Birth
        - Date of Birth
        - Social Security Number
    * Object Identifier
        - Registration Authority
        - Object Type
        - Name
    * Object Protection (Security Label)
        - Registration Authority
        - Object Type
        - CIA
        - Compartment
    * Subject Authorization (Privilege Attribute Certificate)
        - Registration Authority
        - Subject Type
        - Name
        - Authorization(s)

NIST Computer Security Object Register

- Goals
  * National/Federal Registration Authority
  * Unique Name for Service Negotiation
  * Catalogue for Users
  * Information Distribution for Vendors
- Status
  * Draft Rules for Registration
  * NIST/NCSL Support for Operation
  * Request for Registration
  * Seeking National Recognition/Approval
- Registered Objects (Tentative Examples)
  * Other Registration Authorities
  * Cryptographic Algorithms
  * Key Management Systems
  * Security Domains
  * Security Labels

NIST Labeling Approach

- Hold Workshop on Security Labels
- Create Generic Security Label Format
- Test Generic Format with Existing Labels
    * IPSO (Revised)
    * ?
- Synthesize Labels for Hypothetical Domains
- Propose FIPS on Security Labels
    * Separate Standard and/or
    * Incorporate in GOSIP
- Register Domains and Labels
- Develop FIPS on Trusted Distributed System Supporting Multiple Domains and Labels

"DoE Proposal for Security Labeling in GOSIP", C. Douglas Brown (Sandia National Laboratories)

DOE Proposal for Security Labeling in GOSIP

C. Douglas Brown, 2636
Sandia National Laboratories
Albuquerque, NM 87185

I. Introduction and History

Early in 1988, the Department of Energy chartered a Security Labeling Standards Working Group to draft DOE standards for security labeling of network communications. This working group was headed by Dale Sparks of Los Alamos National Lab and included site security managers and technical computer security personnel from the major DOE laboratories and production agencies. A technical subcommittee consisting of Dave Wiltzius, Lawrence Livermore National Lab, Steve Turpin, Los Alamos National Lab, and Doug Brown, Sandia National Labs, was formed to draft the actual proposal for a security labeling standard, with considerable input and review by the Working Group. In January, 1989, a final draft was prepared by the subcommittee and accepted by the Working Group.

The original proposal for a DOE security labeling standard was oriented toward the TCP/IP protocols and was based upon the Revised IP Security Option defined by Captain Michael St.Johns, USAF, in RFC 1038. After this draft proposal for IP security was accepted by the DOE working group, it was modified to fit within the framework of GOSIP Version 1 Connectionless Network Protocol (CLNP) and was reissued later in 1989 under the title "Draft for DOE Use of CLNP Security Options".

II. Use of the Revised IP Security Option.

The DOE proposal adapts and extends the Revised IP Security Option defined by RFC 1038 in the following ways:

1. The Basic Security Option is used as is, and is required on each IP datagram. This option defines the four basic security levels: Unclassified, Confidential, Secret, and Top Secret, with an additional four security level numbers defined as "Reserved for Future Use".

2. An additional Protection Authority Code was requested by DOE and assigned by DCA. The DOE code is 4.

3. The Extended Security Option is not required on each IP datagram. It contains an Additional Security Information field whose contents are undefined by RFC 1038. If the DOE Protection Authority Flag is set, the DOE standard further defines this field to contain the following:

   a. The DOE label version number (currently 1).

   b. An octet reserved for future use.

   c. Two octets containing a Category bit mask.

   d. Up to 14 octets containing a Handling Instruction bit mask. All known DOE handling instructions and caveats are defined in the first 6 octets of this field. The last 8 octets are reserved for site-specific use.

Also, products that support the RFC 1038 Basic Security Option are beginning to appear in the marketplace, including a router that filters IP datagrams based upon the security level contained in the Basic Security Option. Such a router could be used in a network using DOE security labeling, as long as it simply passed the Extended Security Option along unchanged, and the enforcement of access by security categories and compartments could be left up to the hosts.

A third reason to choose this approach is that the DNSIX interface specifications being developed by The Mitre Corporation for the Defense Intelligence Agency specifies the use of the Basic Security Option to communicate security levels and the Extended Security Option to communicate security compartments. In fact, the DNSIX definition for the Extended Security Option is quite similar to the DOE definition, though not identical. In fact, if the DOE labeling standard were modified to reverse the order of the DOE label version octet and the reserved (unused) octet, then the unused octet would appear in the same position within the Extended Security Option as the DNSIX SOURCE field, which is used to qualify the definition of the compartment designator bits following it. If DIA were willing to assign a value of the DNSIX SOURCE field to DOE for its use, the DOE and DNSIX labels could be made compatible.

V. Adaption of the DOE Labeling Standard to GOSIP

The GOSIP Version 1 specification first appeared in June, 1988, and contained a chapter on "Security Options" (Chapter 6). This chapter defined a CLNP security parameter as follows:

1. Parameter Code: 1100 0101

2. Parameter Length: variable.

3. Parameter Value as follows:

   a. ISO Security Format Code

   b. Security Option Type
      Basic Security Option (1000 0010), or
      Extended Security Option (1000 0101)

   c. Security Option Value

The Security Option Types and Values were defined by GOSIP 1.0 to be identical in every respect to the Basic and Extended Security Options as defined by RFC 1038. The DOE Security Labeling Standard as applied to GOSIP fits within the above framework and should be a workable mechanism for performing security labeling in OSI networks both within and between DOE sites.

III. The DOE Security Model

The Basic Security Option is required on each datagram, and if the data being transmitted falls within any of the DOE defined Categories or Handling Instructions, then an Extended Security Option is required as well. These Security Options contain orthogonal components of the security label. The Basic Security Option contains the security level and the Extended Security Option contains the security categories and handling instructions, which are collocated so that they can be treated as a single bit mask, if so desired. (In fact, a number of B1 systems currently available or under development represent security categories and compartments internally as a bit mask of 32, 64, or 128 bits. This would map well into the DOE label format and would permit a simple, efficient implementation of network security labeling.)

The Basic and Extended Security Options are to be used by host systems and trusted intermediary systems (routers) for accepting or rejecting a datagram based upon its security level, categories, and handling instructions. Each host will have an associated accredited security classification range, which is composed of a minimum and maximum security level, a minimum and maximum category bit mask (explained below) and a minimum and maximum handling instructions bit mask. The security classification of each incoming datagram must fall within the range for a host, or that host must reject the datagram following the prescribed out-of-range procedure. In addition, each network interface may be configured with a security classification range. In that case each incoming datagram must fall within the range of the respective interface, or the host must reject the datagram following the out-of-range procedure. Each outgoing datagram must fall within the range of the respective interface, or the datagram must not be sent and the process attempting to send the datagram must be returned an error.

A minimum bit mask represents the set of bits that all acceptable datagrams must contain. A maximum bit mask represents all the allowable bits that may be set in an acceptable datagram. A datagram having any category or handling instructions bits set that are not present in the corresponding maximum bit mask must be rejected.

The default value for the Handling Instructions field is zero. That is, if the Handling Instructions field of a datagram has been omitted (the length field of the Extended Security Option is 5-7), the Handling Instructions are assumed to have a value of zero (no bits set). If the Extended Security Option is omitted, then the Categories are assumed to have a value of zero.

IV. Justification for Use of Basic and Extended Security Options

While a security level could have been incorporated into the Extended Security Option and the Basic Option could have been omitted, it was felt that the use of both options in conjunction with each other was better, as that would maintain compatibility with the original intent of RFC 1038.

DOE Proposal for Security Labeling in GOSIP

C. Douglas Brown

Division 2636

Sandia National Laboratories

Albuquerque, New Mexico

Sandia National Laboratories

CDB:2636:NIST_Gosip:5/26/90

History of DOE Security Labeling Standards

04/88 - DOE chartered Security Labeling Standards Working Group

09/88 - Initial draft prepared for IP

01/89 - Final "Draft for DOE Use of IP Security Options"

06/89 - Modified to fit GOSIP 1.0 CLNP Security Parameter

??/89 - Updated for GOSIP 2.0

Sandia National Laboratories

CDB:2636:NIST_Gosip:5/26/90

DOE Approach to IP Security Labeling

Use RFC 1038 Revised IP Security Option (RIPSO)

Use RIPSO Basic Security Option as is

- Only four security levels (U,C,S,TS)
- Get Protection Authority Flag for DOE

Define Extended Security Option further for DOE

- DOE label version number (1)
- Octet reserved for future use.
- Two octets containing a Category bit mask
- Up to 14 octets containing a Handling Instruction bit mask
    - First 6 octets defined DOE-wide
    - Last 8 octets for site-specific use

Sandia National Laboratories

CDB.2636:NIST_Gosip:5/26/90

The DOE Security Model

Basic Security Option required on each datagram

Extended Security Option optional

- If missing, zeroes are assumed
- Categories and Handling Instructions may be treated as a single mask

Each host has accredited classification range

- Min/Max security levels
- Min/Max security categories
- Min/Max handling instructions

Each network interface on a host may have a classification range

Sandia National Laboratories

CDB:2636:NIST_Gosip:5/26/90

Adapting the DOE Labeling Standard to GOSIP

GOSIP 1.0 Spec Defines a Security Parameter

- Parameter Code: 1100 0101
- Parameter Length: variable.
- Parameter Value as follows:

  ISO Security Format Code

  Security Option Type (Basic or Extended)

  Security Option Value

DOE Basic & Extended Security Options fit GOSIP framework

Multiple Extended Security Options permitted

- Needed by DNSIX but not DOE

Sandia National Laboratories

CDB.2636:NIST_Gosip:5/26/90

Comparison to DNSIX IP Labeling

Both use RIPSO

Both require Basic Security Option

Both define Extended Security Option

- \- Optional for both
- \- DNSIX allows multiple instances
- \- DOE defines an extra unused octet
- \- Both define compartments (categories) as a bit mask

Sandia National Laboratories

CDB:2636:NIST_Gosip:5/26/90

Commercial IP Security Option (CIPSO)

New IP Option Type requested by TSIG (Trusted Systems Interoperability Group)

Domain of Interpretation (4 octets)

- Assigned by registering authority to community of users
- Defines interpretation of security information, e.g., category bit mask

Token ID or Type

- Assigned by registering authority
- Defines format of security info.
- Type 1
  - Security Level (1 octet)
  - Security Categories (8 octets), treated as a bit mask
- Type 2
  - Security Level (1 octet)
  - Security Categories (16 octets)

Sandia National Laboratories

CDB:2636:NIST_Gosip:5/26/90

Similarities in DOE, DNSIX, & CIPSO

Use a bit mask to represent categories and compartments

- (2-16 octets)

Use a field to define interpretation of a category bit mask

- DOE version number (1 octet)
- DNSIX Source field (1 octet)
- CIPSO Domain of Interpretation (4 octets)

Allow different formats of security information

- DNSIX uses Source field
- CIPSO uses Token ID

Sandia National Laboratories

CDB.2636 NIST_Gosip 5/26/90

Differences in DOE, DNSIX, & CIPSO

DOE & DNSIX use Basic Security Option

- Only 4/8 security levels defined
- Levels have decreasing values (must map into O/S dependent values)
- A few implementations exist

CIPSO incorporates security level with categories in Token IDs 1 and 2

- 256 possible levels
- Map well into MLS operating systems
- No implementations exist

Sandia National Laboratories

CDB.2636 NIST_Gosip 5/26/90

Principles for Defining Security Labels

Must have a Domain of Interpretation field

- Qualifies interpretation of bits
- Not administered exclusively by Defense agencies
  - To be responsive to commercial sector
  - Perhaps shared administration

Need several Subtypes (Token IDs)

- To distinguish various types/formats
- To allow extensibility
- Need at least a category bit mask (8-16 octets) and security level octet

Prefer only ONE label type

- To encourage implementation by vendors

Sandia National Laboratories

CDB:2636:NIST_Gosip:5/26/90

Summary and Conclusions

<u>DOE needs are met by RIPSO & GOSIP 1.0</u>

- 4 levels & category bit mask are enough
- Could be made DNSIX-compatible

<u>Will have both RIPSO and CIPSO label formats in the IP world</u>

<u>Should attempt to merge to one label format in the GOSIP world</u>

Sandia National Laboratories

CDB.2636:NIST_Gosip.5/26/90

Proposed DOE Category Bits

0 NONE Categorized as "No category"
1 SU Sensitive Unclassified
2 UCNI Uncl. Controlled Nuclear Info.
3 PARD Protect As Restricted Data
4 NSI National Security Information
5 FRD Formerly Restricted Data
6 RD Restricted Data
8 SCI Sensitive Compartmented Info.
(Qualified by one or more bits in the site dependent area)

Sandia National Laboratories

CDB:2636:NIST_Gosip:5/26/90

Presentation Slides, David Chizmadia (National Computer Security Center

PURPOSE

- Common Understanding
- Community Needs
- Pros/Cons of Possible Implementations

COMMON UNDERSTANDING

- Role of Security Policies
- Marking vs Labeling
- Label Design Goals
- Choosing Subjects & Objects
- TCSEC Requirements

COMMUNITY NEEDS

- Models
- Design Considerations
- ?

IMPLEMENTATIONS

- Hierarchical Levels
- Non-Hierarchical Levels
- Non-Secrecy Labels

"The Secure Network Password", A.A. Eshgh, P.H. Wiedemann (SSDS, Inc.)

The Secure Network Passport

P.H. Wiedemann and A.A. Eshgh
SSDS, Inc.

As any network grows, there comes a point when centralized security administration is no longer practical. Decentralized (distributed) security administration requires a mechanism and an agreed upon set of standards to be useful and effective. The two principal components of such mechanism are privilege (credentials) security and data transport security.

As its title implies, the first part of this paper modestly proposes that the privilege security problem is similar to the issues of credentials used by envoys in international diplomacy and by ordinary citizens as they travel among countries that do not share a centralized administration. The second part of this paper proposes a scheme of security labeling for Protocol Data Units, which constitutes the principle data transport element for the exchange of data in modern networks.

Part 1 - The Network Passport

In this paper, many references will be made to Protocol Data Unit (PDU). PDU is a term used by the International Standards Organization to describe the data communications equivalent of an ordinary postal letter. A letter contains two primary constituents, the envelope and the contents of the envelope. Likewise, a PDU consists of two primary constituents, the header portion and the data portion. Like the envelope, the header usually includes destination (recipient) address and the source (sender) address. A PDU header also contains validation information and other miscellaneous delivery control information. Security handling instructions, where used, also are found in the header. The data are analogous to the content of the envelope. Just as proper mail delivery service concerns itself only with the instructions on the envelope, a network only examines the header during the execution of its delivery task. Just as envelopes may be placed inside other envelopes (say, a personal letter inside an express delivery envelope), PDUs commonly become the data portion of other PDUs, and may be nested several layers deep, each layer fulfilling its own role in the end to end information transfer function across a network. Some examples of PDUs include packets, frames, and datagrams.

The Basic Elements and Analogies

A security domain consists of a realm over which a single security administrator has control. This may be a host, a communications network, a subnet, etc. It is analogous to an independent state or country.

Each security domain may establish a level of trust for every security domain which one of its own users intends to access. It is analogous to the establishment of diplomatic relations among countries. The level of trust is based on the home domain's opinion of how well security is enforced in each of the "foreign" domains. An independent third party (clearinghouse), whose opinions in these matters are trusted by both domains may be used as a mediating entity and has many advantages. Such clearinghouse will usually set standards for, and examine the adequacy of, the security environment of a candidate domain and apply a security rating based on the opinion of the examiners. This is analogous to an organization such as the European Economic Community (EEC) in matters of trade and tariffs.

Each security domain is responsible for identifying and authenticating each of its users internal to that domain and for establishing a portfolio of local permissions or privileges (security clearances, in military and intelligence systems) for that user. This is analogous to a national information file or "dossier" on each citizen.

By convention with the other domains, a unique and unforgeable token is assigned by each domain to each user. Here again, a third party, trusted by all participating entities, is very helpful as a neutral catalyst to foster agreement on form and content. The token is trusted, to the same degree as the domain is trusted by other domains, to be intimately associated with the user's complete credentials set (dossier) located at the user's home domain. This is analogous to a passport. The token's form, content, coding and such is standardized by agreement among domains or through the third party. This token becomes the user **Identification security label** that is attached to each PDU representing a user request for service from another domain.

The token may contain embedded and unforgeable information attesting that a specific foreign domain has inspected and pre-authorized a user for access to it's domain. This is analogous to a visa attached to the passport. One or more such visas (one for each foreign domain) may be associated with the passport. Depending on the agreement (diplomatic relation) between any particular pair of domains, visas may or may not be used.

Presence of a passport provides only a user's ID and authentication, not permissions/clearance, therefore simplifies access control because access privilege information is retained by target, not source of the access request. Use of the passport allows reliable association between the accessing entity and the set of privileges maintained by (and likely differing for) each access target.

The Passport in Operation.

Because the passport contains, in its simplest form, only identification and authentication information about the user who wants service from the foreign domain, the format can be very simple. Unlike more common security labels, which contain privilege/clearance information, the format and meaning are easy to discern and can be honored regardless of the quality of security enforced by the "target" system.

Integrity seals are also required with the passport to ensure that the access or service requested was in fact requested by the user and not inappropriately requested on the users behalf by an unauthorized agent somewhere along the PDU's journey. Authorized and trusted agents may be able to request vicarious accesses and services on the user's behalf.

Typically, a user would have pre-established, to the satisfaction of the foreign system, the need and authority to access information in, and to receive data from, the foreign domain. When a Service Request PDU arrives at a domain checkpoint (analogous to a check at a border or a building or an office) the immediate target domain checks the user's passport against locally held privilege lists to determine what the user may do while in that domain. If a visa is attached, it verifies that the seal is intact and then authorizes the user to proceed without reference to the privilege table because the visa is evidence that the privilege has been pre-established.

If the user has not pre-established need and authority for service/access at a foreign domain, that domain can "detain" the PDU and send a request to the domain of origin, obtainable from the passport (the same way real passports are identified with the country of issuance). The user or the security administrator, or both, may (and usually will) place limitations on the kind of data accessible in a user's dossier by the foreign domain, based on that domain's need to know.

Advantages

• The passport can be distributively administered

• The simplest form can be quickly established because it only contains ID and authentication information, requiring no lengthy and controversial negotiations on meanings of privileges. Most countries in the world agree, regardless of their politics or ideologies, on use of passports.

• Use of visas allow privileges to be included in passports to simplify border decisions.

• The token should be small enough to be included as a label in every PDU that directly or vicariously represents a user.

• The privilege information can be locally held at destinations of service/access requests instead of at origins, simplifying administration at each domain and reducing trust negotiations among domains.

• The same credentials (passport, token) can be used at convenient intervening checkpoints (security interfaces between major components) such as between a workstation application and its communications system (airport passport check), at the entrance to and exit from a network, at the entrance to a host or other network, at the application within a host or server. This reduces the amount of credential information carried by the PDUs and increases useable communication bandwidth.

• The passport applies the benefits of an already-working solution that has stood the test of time in world diplomacy to distributed systems security.

Part 2 - A Network Security Labeling Scheme

Once the communications relationship has been established through the use of passport, PDUs containing data may then be interchanged. Such PDUs may also traverse multiple domains between their source and destination. Since data of differing classifications will likely pass through the same interfaces and in many cases share media, a system must be provided that

• Identifies the classification of each PDU, and

• Separates PDUs such that during transit from source to destination the data of one PDU cannot become mixed or interchanged with the data of another PDU.

Several techniques are used for both purposes and include the following:

Spatial - Different medium for each classification

Temporal - Shared medium with allocated time slots for each classification. It essentially uses Time Division Multiple Access (TDMA) techniques for purposes of security enforcement.

Spectral - Shared medium with allocated frequency bands in the radio or light spectrum allocated for each classification. It uses Frequency Division Multiple Access (FDMA) techniques for purposes of security enforcement.

Cryptoral - Shared medium where the data is purposely scrambled to obscure its meaning, where the the scrambling algorithm and keys of the sender, different for each classification, are known only to authorized receivers. This technique includes the use of non-digital mechanisms used for the same purpose, such as spread spectrum technology in the radio or light spectrum.

Labels - Shared medium where digital security labels are intimately associated (typically applied contiguously in time or in storage areas) with the data to which they apply. In networks using PDUs, each PDU to be handled securely carries such a label.

The remaining discussion will concern itself only with the use of labels. This paper also does not address labels used for the purpose of providing integrity for PDUs.

The use of security labels, as a means of distinguishing classifications of PDUs in networks and to separate their data, requires additional protections to constitute a complete protection system. As such,

- The label itself must be protected against alteration (label integrity protection). Techniques such as encrypted checksums or other encrypted integrity characters have been found useful.

- Because the presence of labels do not preclude reading of the labels of any PDU, the labels themselves may need be protected against being read by unauthorized readers (label secrecy protection). This could be provided through the use of encryption or through physical protection of the media.

- The PDU's data portion, which is associated with the label, may also need to be protected against both integrity and secrecy attacks. This could again be provided through use of encryption or through physical protection of the media.

- The entire PDU must be protected against Denial of Service attacks.

Families of Classifications

Two classification label methods are used today - hierarchical and non-hierarchical. For each system, access controls can be further divided into two broad categories - mandatory and discretionary. Each will be described below.

Hierarchical Classification

This method classifies data into one of several contiguous, hierarchical, discrete levels. Any PDU must be classified into one and only one level. A label reflecting the level requires only a single value to represent the degree in which its associated data has been classified. An example of such method is NATO's levels that include Confidential, Cosmic, etc.

Non-Hierarchical Classification

This method classifies data into one or more classification categories, all of which are on the same hierarchical level. Different categories of such classification therefore lack the higher/lower relationship associated with hierarchical classification categories. Because of this, data in a PDU may simultaneously be classified into more than one category and perhaps into all categories of the specific classification system. Examples of classification systems include DoD, DIA, NSA, CIA, NATO, DoE, Banking Industry, etc. A label reflecting multiple categories of non-hierarchical classification must therefore be able to simultaneously indicate several, and perhaps many different values, each representing a specific classification category. Each value represents a discrete clearance from among a set of clearances ranging from a few to thousands, with many additions and deletions of discrete values from time to time within a classification system. Various compartmental security categories as defined by DoD are examples of this classification method.

Each non-hierarchical clearance value can be represented in Boolean form as true or false. For any specific non-hierarchical classification, these Boolean values can be represented by a single bit in the security label. It is suggested that the most commonly used convention be adopted, which equates a binary one (1) to a Boolean True condition and a binary zero (0) for a Boolean False condition.

Mandatory Controls

Mandatory security classification enforcement applies when certain security protections of the PDU are required by regulation or law and not within the discretion of the owner of the information to apply or ignore. Where labels are used, they must be honored by all components of the network through which the PDU passes. Labels indicating classifications that reflect mandatory controls must be distinguished from those indicating classifications that reflect discretionary access controls.

Discretionary Controls

Discretionary security classification enforcement applies when certain security protections of the PDU are imposed on a need-to-know basis and may be changed by security administrators based on their best judgement. Where labels are used, they may or may not be honored, depending on whether the network component enforces discretionary access controls or not. Labels indicating classifications that reflect discretionary controls must be distinguished from those indicating classifications that reflect mandatory access controls.

Combinations of Classification Families

Since either or both hierarchical and non-hierarchical labels may be subjected to either or both mandatory and discretionary controls, four combinations of classification families must be treated by the labeling system to be used. The following system is proposed to effectively deal with all four combinations.

- Hierarchical Mandatory

- Non-Hierarchical Mandatory

- Hierarchical Discretionary

- Non-Hierarchical Discretionary

Each combination needs separate sublabel when applied to a PDU. The following scheme is modestly proposed for consideration.

The minimum security label consists of one octet as shown in Figure 1. If all bits

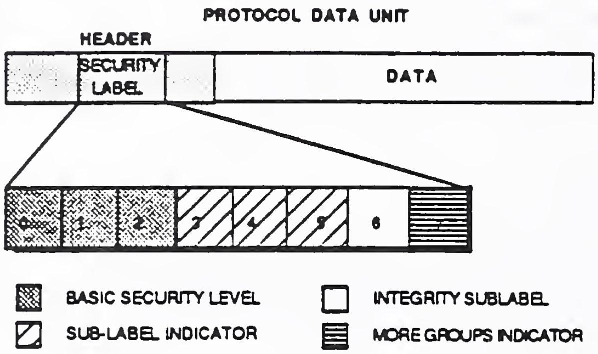

FIGURE: 1
MINIMUM SECURITY LABEL

in this octet are zeros (0) it indicates that there is no security label for that PDU. If any bit in this octet is not zero (0) then the octet is decoded as follows:

> The first three bits (indicated as bits 0, 1, and 2 here) are combined to indicate one of eight possible **Hierarchical, Mandatory** security levels. Because there is only one system of mandatory, hierarchical classification, only a single value is required. The proposed coding scheme for this security sub-level is presented blow and is depicted in Table 1.
>
> > 0 = Unclassified. It is assumed that no level lower than unclassified will be established.
> >
> > 2 = FOUO or Sensetive but Unclassified. This leaves room on either side of this level to introduce a new hierarchical level. Since there have been several recent developments in the "sensitive but unclassified" arena, additional levels in this region are more likely than for other levels.
> >
> > 4 = Confidential. This and the next two levels have been stable for a considerable number of years and have every appearance of continuing to do so in the foreseeable future. For this reason, this proposal suggests no gaps between these levels for future creation of new intervening levels.
> >
> > 5 = Secret
> >
> > 6 = Top Secret. One more "slot" (value 7) is left above the Top Secret level to allow for future creation of one more level above Top

Secret. Since no such new level has appeared in a number of years, the allocation of one slot only would seem to be sufficient.

TABLE: 1
.MANDATORY, HIERARCHICAL SECURITY SUB-LEVEL CODING SCHEME

VALUE	DESCRIPTION
0	UNCLASSIFIED
1	RESERVED FOR EXPANSION
2	FOUO
3	RESERVED FOR EXPANSION
4	CONFIDENTIAL
5	SECRET
6	TOP SECRET
7	RESERVED FOR EXPANSION

The next three bits (indicated as bits 3, 4, and 5 here) indicate whether (1) or not (0) sublabels for each of the three remaining combinations of classification families below are present as part of this security label. The sublabels, if used, must appear in the agreed upon order. One such order is suggested below. A value of zero (0) in all three bits indicates that this PDU has no further sublabels and that the entire security label consists of but a single octet.
Bit number six may be utilized to enhance the label integrity. Once it is set (1), it could be used to indicate the presence of an integrity sublabel.

The last bit indicates if this is the only group in the first sublabel. A (1) value suggests the presence of additional group(s) in the Mandatory, Hierarchical sublabel. Conversely, a (0) value indicates no other group is present in the first sublabel.

If multiple systems of Mandatory, Hierarchical classifications are to be accommodated, a structure similar to that shown below under Hierarchical, Discretionary classifications may be used instead.

Non-Hierarchical, Mandatory Sublabel

Because:

• There are several classification systems for non-Hierarchical, Mandatory classification with no deterministic mapping of classifications between systems (for example, two secret level labels associated with DIA and NATO or even within the same classification system may not map), and

• Some of those systems have a large number of discrete values and

• A single PDU may contain any combination of those discrete values

The following sublabel scheme is proposed to represent this combination of classification families:

- The sublabel is of variable length but is divided into discrete groups.
- The number of groups is not restricted by this proposal but may have external restrictions such as those that limit total PDU length.
- The groups fall into two types, bit-mapped and discrete. Figure 2 illustrates the proposed sublabel group scheme.

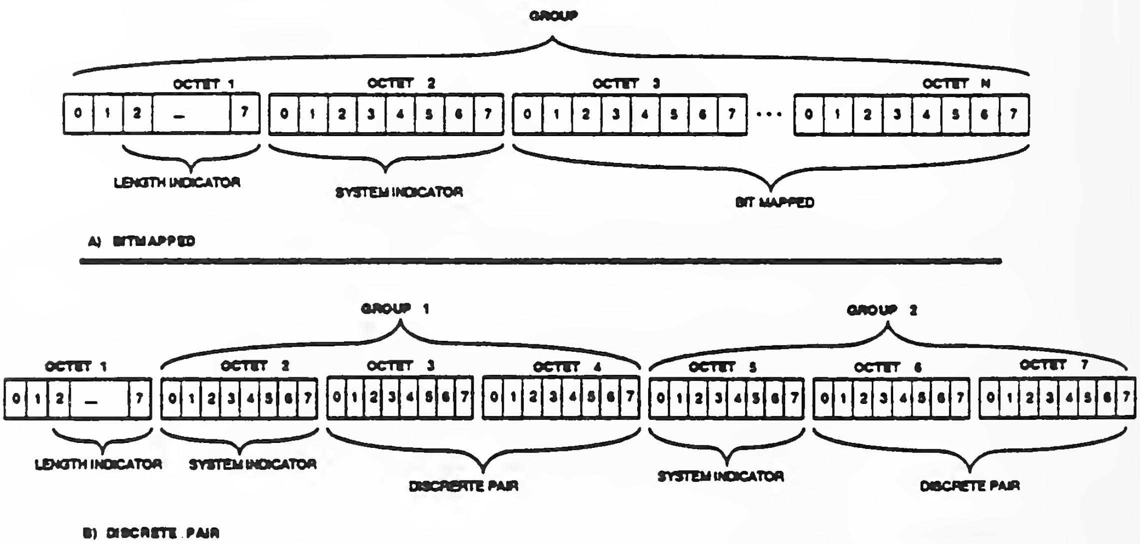

FIGURE 2
NON-HIERARCHICAL MANDATORY SUBLABEL GROUP SCHEME

A bitmapped group consists of two parts. The first part indicates which one of a number of specific (registered) classification systems is represented by this group (meaning that one group may not represent more than one classification system. The second part is a bit map wherein each bit indicates whether (1) or not (0) the PDU is classified with a specific non-hierarchical classification belonging to the classification system identified by the first part of this group.

A discrete-pair group. Such group consists of pairs of indicators. The first indicator in each pair indicates one specific (registered) classification system. The second indicator indicates a specific non-hierarchical classification from within the system indicated by the first indicator.

The first Octet of this sublabel is coded/decoded as follows and shown in Table 2:

The first bit indicates whether (1) this group is a bitmapped group or (0) a discrete-pair group.

The second bit indicates whether (1) or not (0) more groups (of the same sublabel) follow this group.

The six remaining bits are combined to indicate the length, in octets, of the remainder of the group. This allows groups of from two to 65 octets in length.

TABLE: 2

NON-HIERARCHICAL MANDATORY SUBLABEL FIRST OCTET DESCRIPTION

BIT	POSSIBLE VALUES	DESCRIPTION
0	0	DISCRETE PAIR GROUP INDICATOR
	1	BITMAP GROUP INDICATOR
1	0	NO OTHER GROUPS OF THIS SUBLABEL WILL FOLLOW THIS GROUP
	1	ANOTHER GROUP OF THE SAME SUBLABEL WILL FOLLOW THIS GROUP
2 THROUGH 7	0 AND 1	GROUP LENGTH INDICATOR

The second octet is encoded/decoded as follows:

If the group is a bitmapped group, the second octet indicates the classification system (one of 256). The remaining octets in this group each indicate whether (1) or not (0) the PDU is classified with a specific non-hierarchical classification belonging to the classification system identified by the second octet in the group.

If the group is a discrete-pair group, then the second octet indicates the classification system (one of 256) while the following two octets indicate the specific non-hierarchical classification belonging to the classification system identified by the second octet in the group (one of 64,536). This three-octet pattern continues for the group length indicated in the group's first octet.

This arrangement reduces the number of octets required, both accommodating environments such as DoD where there are many classification systems with few classes applied to the PDU as well as environments such as CIA where there are few, or one, classification system(s) with many classes applied to the PDU.

Hierarchical Discretionary Separation

This arrangement is similar to the hierarchical, mandatory scheme in that only one value per classification system is required, but there may be several classification systems. Also, the number of levels may be larger than the eight (8) level provided for the hierarchical, mandatory combination as is the case with some banking environments. The following encoding/decoding scheme is suggested for this sublabel:

The first octet in this sublabel is coded as follows and presented in Table 3:

The first bit is unused

The second bit indicates whether (1) or not (0) more groups (of the same sublabel) follows this group.

The six remaining bits are combined to indicate the length, in octets, of the remainder of the group. This allows groups of from two to 65 octets in length.

The remaining octets of this sublabel consists of pairs of octets. The first octet in each pair indicates one of 256 classification systems while the second octet indicates one of 256 hierarchical security levels.

TABLE 3
HIERARCHICAL DISCRETIONARY FIRST OCTET SCHEME

BIT	POSSIBLE VALUES	DESCRIPTION
0	0 AND 1	NOT USED
1	0	NO OTHER GROUP OF THIS SUBLABEL WILL FOLLOW THIS GROUP
	1	ANOTHER GROUP OF THIS SUBLABEL WILL FOLLOW THIS GROUP
2 - 7	0 AND 1	GROUP LENGTH INDICATOR

Non-Hierarchical Discretionary Separation

The coding scheme of this combination is identical to the non-hierarchical, mandatory scheme. Such a scheme can fully support the security needs of organizations such as banking communities where an administrator can decide which office or department may have access to which information.

Advantages

The proposed approach offers a number advantages. It:

- Conserves the number of protocol octets when used in simple security environments. For example, only a single octet is needed if only the levels from unclassified through Top Secret are to be applied to PDUs.

- Accommodates complex security environments where a PDU may belong to several classification systems and use both hierarchical and non-hierarchical classifications under both mandatory and discretionary security policies.

- Supports variable length security labels

- Provides complete access control identification and separation security for PDUs

- Allows for both a variety of classification systems as well as a variety of hierarchical and non-hierarchical classifications within each system.

- May be applied to both Government and non-Government environments.

"Labelling Communicated Data In OSI", C. Gray Girling (Top Express Ltd., U.K.)

Labelling Communicated Data in OSI

C Gray Girling — 25 April 1990

Abstract

This document establishes a number of options for a security labelling strategy based on standard Open System Interconnection (OSI) protocols and services.

Contents

TXSG-90-CGG-23

1. Introduction

This document provides some investigation of the means to provide the access control service, given a specific requirement: that access control should be providing using security labels at the application layer. This document investigates the options visible *within an endsystem* for the use of security labels in the implementation of an OSI communications protocol stack. Its scope is broader than that implied by the term Open System Interconnection because that would normally address issues visible only *outside an endsystem.* This document categorizes the strategies *within an endsystem* which can be used to support labelling.

Although access control is an application layer requirement, the problem is examined from the point of view of a general layer (referred to as the (N)-LAYER). The interaction between solutions chosen for separate layers is then constrained so that the combination of layer solutions can support access control at the application layer.

2. Access Control Requirements

It is required that access control be exercised on the basis of a set of partially ordered SECURITY LABELS. The access control mechanism is to associate security labels with transmitted data, and ranges of security labels with application-process-invocations (APIs). It must ensure *at least* that no data associated with a security label outside the intersection of the ranges associated with the two participating APIs is transmitted from one to another.

3. Labelling Requirements

The top of the OSI protocol stack provides the ability to send and receive a number of application-protocol-data-units (A-PDUs) which may used in either a connection orientated mode (as part of an application association) or in a connectionless mode (as part of an application-unitdata service).

A stream of communicated data can be divided into sections visibly associated with one or a set of security labels each called a LABELLED FIELD. The same stream of data can also be divided into A-PDUs. In theory these two methods of subdividing the stream of data are independent, however, there are advantages in synchronizing them. For an application association the following options are considered:

- all A-PDUs may be associated with the same label or set of labels;
- each A-PDU may be associated with a different label or set of labels; or,
- each A-PDU may be divided into parts each of which are associated with a different label or set of labels.

That is, an association labelled field may encompass the entire association, one A-PDU or only one of many parts of an A-PDU.

Thus, labelling requirements are relevant at the *application layer* (and only indirectly to lower layers). They can be divided, usefully, into six classes:

(1) FIXED LABEL ASSOCIATION: the requirement that an application association should be able to be created all of whose A-PDUs are visibly associated with a *single* security label (i.e. labelled field = the whole association);

(2) FIXED LABEL-SET ASSOCIATION: the requirement that an application association should be able to be created all of whose A-PDUs are visibly associated with a *single* set of security labels (i.e. labelled field = the whole association);

(3) VARIABLE LABEL ASSOCIATION: the requirement that an application association should be able to be created each A-PDUs of which are visibly associated with a *single* security label, but not necessarily the same one (i.e. labelled field = one A-PDU);

(4) VARIABLE LABEL-SET ASSOCIATION: the requirement that an application association should be able to be created each A-PDUs of which are visibly associated with a *single* set of security labels, but not necessarily the same one (i.e. labelled field = one A-PDU);

(5) MULTIPLE LABEL ASSOCIATION: the requirement that an application association should be able to be created each A-PDU of which is divided into a number of fields each visibly associated with one of *many* security labels (i.e. labelled field = a fraction of an A-PDU); and,

(6) MULTIPLE LABEL-SET ASSOCIATION: the requirement that an application association should be able to be created each A-PDU of which is divided into a number of fields each visibly associated with one of *many* sets of security labels (i.e. labelled field = a fraction of an A-PDU).

4. The Labelling Problem

The position in the communications architecture at which the access control service is provided may be at a layer below the application layer.[1] Thus there is a gap between the position at which access control is required (the application layer) and the position at which it is provided. This latter layer (or sublayer), at which data labelling is manifest in the layer protocol and service, will be referred to as the (L)-LAYER.

---

1 As, for example, it is in the US SDNS SP3 (network layer) and SP4 (transport layer) sets of protocol.

This document provides a review of a number of ways in which the labelling requirement at the application layer can be mapped on the label provision at the (L)-layer. This involves the specification of a selection of local mechanisms applicable to each intermediate layer. Since the following is generic to each of the intermediate layers the analysis refers to the (N)-layer.

5. Labelling Facilities

OSI distinguishes, with respect to the (N)-layer, the units of data that the layer service transfers on behalf of its user ((N)-service-data-units, or (N)-SDUs) and the coded units of data that the layer protocol actually uses to accomplish the transfer ((N)-protocol-data-units, or (N)-PDUs).

For the purposes of this document the part of the OSI layer (N)-service that associates security labels with data carried will be called a LABELLING (N)-FACILITY. This is divided into six main types according to how a labelled field (called an (N)-LABELLED-FIELD in this context) maps on to an (N)-SDU and whether a single label or a set of labels are indicated:

- FIXED-LABELLING (N)-FACILITY

  [FIXED-SET-LABELLING (N)-FACILITY]

  an (N)-facility that associates the same label [set of labels] with all of the data of each each (N)-SDU in an (N)-connection (i.e. the (N)-labelled-field is the entire (N)-connection);

- VARIABLE-LABELLING (N)-FACILITY

  [VARIABLE-SET-LABELLING (N)-FACILITY]

  an (N)-facility that associates a label [set of labels] with all of the data of each each (N)-SDU in an (N)-connection (but not necessarily all the same – i.e. the (N)-labelled-field is an individual (N)-SDU); and,

- MULTI-LABELLING (N)-FACILITY

  [MULTI-SET-LABELLING (N)-FACILITY]

  an (N)-facility that can associate different labels [sets of labels] with different fields in each (N)-SDU of an (N)-connection (i.e. the (N)-labelled-field is a part of an (N)-SDU).

The nomenclature X-(SET-)LABELLING is used to denote either X-labelling or X-set-labelling.

In general, mechanisms supporting fixed-(set-)labelling will be simpler than those supporting variable-(set-)labelling, which will in turn will be simpler than those supporting multi-(set-)labelling. Potential additional local features of the (N)-layer interface to support these types of labelling might be as follows:

- fixed-(set-)labelling –

  either a static association, or at most an additional interface element embedded into a connection establishment request;

- variable-(set-)labelling –

  an additional interface element embedded in each SDU specification; and,

- multi-(set-)labelling –

  additional data structuring defining fields independently to those delimited by SDUs; or additional interface elements within each SDU specification defining subfields each with associated labels (or sets of labels).

Mechanisms supporting labels rather than label sets will tend to be a little simpler since only one label needs to be specified. However, label sets will often be defined using an upper and lower bound in conjunction with some fixed ordering (i.e. as a range) – which is equally simple.

For connectionless mode operation fixed-(set-)labelling and variable-(set-)labelling are not distinguished.

6. Mechanisms to Support Labelling (N)-Facilities

In providing a labelling (N)-facility some means is required to bind security labels to (N)-labelled-fields. This data to label binding can be regarded conversely as the separation of data associated with different security labels. This may be achieved either with or without the support of the local endsystem's Trusted Computing Base (TCB).

With TCB support

With TCB support it may be assumed that data can not exist within the implementation of the protocol stack which is not tightly bound to a security label. Two main kinds of support may be envisaged:

(1) data is held by processes and labelled according to a label (or set of labels) associated with the process; or,

(2) data is labelled according to a capability scheme (possibly with a set of labels) and can be manipulated securely through the capabilities, in a limited way, by

any process.

Inevitably the consequence of such implementations must be that the security label (or set of labels) associated with data presented to any layer is derivable (either bound to a process or to the data). At the (L)-layer the labelling service is provided directly in terms of specific protocol elements. It is thus merely a matter for the protocol implementation of the (L+1)-layer to determine the boundaries of (L+1)-labelled-fields and present each labelled field (or segment of it) with its derived label to the labelling service provided by the (L)-layer. (That is, a multi-class (L+1)-PDU would be split into separately labelled (L)-SDUs, and a single-class one would be copied as a single labelled (L)-SDU.) Thus the way in which data is associated with given labels in the protocol stack need not be considered.

The remainder of this document considers the case in which an operating system (i.e. the TCB) does not automatically keep track of the labels to be associated with data or where it cannot provide the required degree of separation.

Without TCB support

Without TCB support the label associated with individual items of application data must be maintained within the protocol stack. This involves the iterative support of labelling (N)-facilities by labelling (N-1)-facilities in order to support labelling application-facilities by the labelling (L)-facilities.

Mechanisms can be provided which support fixed-, variable- or multi- (set-)labelling[2] (N)-facilities on fixed-, variable- or multi- (set-)labelling (N-1)-facilities, with the exception that set-labelling (N)-facilities cannot be supported using labelling (N-1)-facilities, because there will be no general means to identify a single label associated with an item of (N)-layer data in that case. This section enumerates the mechanisms involved and thus addresses the choices available at each layer.

Note that many of the mechanisms described are purely local to an open system and do not require supporting protocol elements. This case is highly desirable since existing standard protocols and services do not provide these explicit means to support labelling facilities. The requirement for inclusion of supporting protocol elements would therefore make open communication impossible. It must also be noted that the requirement for non-standard local mechanisms (even though they may require no additional protocol support) will result in a lack of suitable "off the shelf" protocol implementations.

---

[2] This nomenclature is an abreviated form for fixed-(set-)labelling, variable-(set-)labelling, or multi-(set-)labelling.

Control protocol element labelling

Whether or not TCB support for data separation is used, each protocol must choose a secure, rational and consistent algorithm for labelling their control elements (such as resets, synchronization points, or management data). These elements are not derived directly from transferred information and so their labelling is not easily determined. There are a number of mechanical choices:

(1) when a label or set of labels are associated with an (N)-connection, use that label or set of labels (if possible);

(2) when a label or set of labels was used in initiating an (N)-connection, use that label or set of labels for all subsequent control elements;

(3) when a set of labels are associated with the (N)-connection use the "highest" or "lowest" in the set; or,

(4) use a fixed label or set of labels.

Each of these options represents a valid solution and no selection from them is made here.

OSI Vocabulary

The following relationships between (N)-connections and (N-1)-connections, as described in [ISO84], are used:

- SEGMENTING

  A function performed by an (N)-entity to map one (N)-SDU onto multiple (N)-PDUs.

- REASSEMBLY

  A function performed by an (N)-entity to map multiple (N)-PDUs onto one (N)-SDU. The reverse function to segmenting.

- BLOCKING

  A function performed by an (N)-entity to map multiple (N)-SDUs onto one (N)-PDU.

- DEBLOCKING

  A function performed by an (N)-entity to identify multiple (N)-SDUs which are contained in one (N)-PDU. It is the reverse function of blocking.

- CONCATENATION

  A function performed by an (N)-entity to map multiple (N)-PDUs onto one (N-1)-SDU.

- SEPARATION

  A function performed by an (N)-entity to identify multiple (N)-PDUs which are contained in one (N-1)-SDU. It is the reverse function of concatenation.

- SPLITTING

  The function within the (N)-layer by which more than one (N-1)-connection is used to support one (N)-connection.

In addition the following term is used to describe the cases distinguished by [ISO84] in the session layer:

- CONNECTION-SEPARATION

  The function within the (N)-layer by which more than one (N-1)-connection is used consecutively to support one (N)-connection.

6.1 Supporting a fixed-(set-)labelling (N)-facility

Fixed-(set-)labelling (N)-facility can be supported either using a fixed-(set-)labelling (N-1)-facility or a variable- or multi- (set-)labelling (N-1)-facility. Each of these options is examined in turn.

6.1.1 Using a fixed-(set-)labelling (N-1)-facility

In the case of a connectionless (N-1)-service each (N-1)-SDU is associated with the single label associated with the fixed-labelling (N)-facility.

The (N)-layer mechanism involved in the case of a connection orientated (N-1)-layer must ensure that at any given time there is a one-to-one correspondence between a (N)-layer connections and the (N-1)-layer connection on which it is based, and that each (N-1)-layer connection is associated with the same single label or set of labels associated with the fixed-(set-)labelling (N)-facility.

Note that the one-to-one correspondence between the (N)- and (N-1)- layer connections may or may not include a one-to-one correspondence between the connection lifetimes. The only recognized case in which (N)- and (N-1)- connections might not have the same lifetime is at the session layer where a transport-connection may support a number of consecutive session-connections (splitting) or a session-connection might be supported by a number of consecutive transport connections (connection-separation).

Should either of these cases be used, mechanism will be required to ensure that, after a change in either the (N)- or (N-1)- connection, the single label, or set of labels associated with the two connections remain in correspondence.

If a fixed-labelling (N)-facility uses a fixed-set-labelling (N-1)-facility the set of labels used by the (N-1) facilities consists of only one element – the one associated with the (N)-facility.

It is not possible, in general, to support a fixed-set-labelling (N)-facility using a fixed-labelling (N-1)-facility.

6.1.2 Using a variable- or multi-(set-)labelling (N-1)-facility

The (N)-layer mechanism involved uses the (N-1)-facility to label each of the labelled fields in the (N-1)-SDU with the fixed label or set of labels associated with the fixed-labelling (N)-facility.

If a fixed-labelling (N)-facility uses a variable- or multi-set-labelling (N-1)-facility the set of labels used by the (N-1) facilities consists of only one element – the one associated with the (N)-facility.

It is not possible, in general, to support a fixed-set-labelling (N)-facility using a variable- or multi-labelling (N-1)-facility.

6.2 Supporting a variable- or multi- (set-)labelling (N)-facility

Variable- or multi- (set-)labelling (N)-facility can be supported either using a fixed-(set-)labelling (N-1)-facility or a variable- or multi- (set-)labelling (N-1)-facility. Each of these options is examined in turn.

6.2.1 Using a fixed-(set-)labelling (N-1)-facility

There are four methods described: set labelling, splitting, connection-separation and splitting & connection-separation together.

Set labelling

This method is applicable only to the use of a fixed-*set*-labelling (N-1)-facility.

A variable- or multi- (set-)labelling (N)-facility can be supported using a fixed-set-labelling (N-1)-facility by finding a fixed set of labels that encompasses those that are to be used on a connection by the (N)-facility and using that set to label each labelled field in the (N-1)-connection.

Note, however, that information regarding the precise label or label set associated with the (N)-connection's labelled field is lost and so the received labelled fields may be associated with a superset of the labels transmitted. The association of received information with a set of labels precludes the support of variable- or multi-labelling (N)-facilities in this way.

Splitting

Where the (N)-protocol supports the splitting of an (N)-connection over a number of concurrent (N-1)-connections it is possible for that layer to perform this splitting on the basis of labels (or sets of labels) associated with labelled fields in the (N)-connection in such a way that each of the (N-1)-connections supports data associated with only a single fixed class (or set of classes). That is, a single-labelling (N-1) facility can be used.

Connection-separation

Similarly where the (N)-protocol supports the connection-separation of an (N)-connection over a number of consecutive (N-1)-connections the separation can be performed on the basis of labels (or sets of labels).

Splitting and connection-separation

Where both splitting and connection-separation are possible a fixed size cache of (N-1) connections can be maintained – each associated with a different label or set of labels – which can be closed and reopened with a new label (or set of labels) when the cache does not contain an appropriate (N-1)-connection.

Splitting is not expected to be a feature of existing layer implementations outside the transport layer. Connection-separation is not expected outside the session layer.

Thus, in each case where splitting is used outside the transport layer, or connection-separation is used outside the session layer special purpose local mechanisms are likely to be required to perform the appropriate re-combination at the destination. These mechanisms will be complicated by the need to render the control elements of a variable- or multi-(set-)labelling (N)-protocol instance correctly in the number of single-labelling (N-1)-protocol instances that operate either consecutively or in parallel to support it.

If a variable- or multi- labelling (N)-facility uses a fixed-set-labelling (N-1)-facility the set of labels used by the (N-1) facilities consists of only one element – the one associated with the (N)-facility.

It is not possible, in general, to support variable- or multi- set-labelling (N)-facility using a fixed-labelling (N-1)-facility.

6.2.2 Using a variable- or multi-(set-)labelling (N-1)-facility

The mechanisms involved here require the use of the same label (or set of labels) used in a labelled field of an (N)-connection to be used for the corresponding labelled field in the (N-1)-connection.

As can be seen from their definitions, the effects of segmentation & reassembly, blocking & deblocking, concatenation each have the effect of destroying the alignment between (N)-SDUs and (N-1)-SDUs (this alignment being established via two sub-mappings: one between (N)-SDU and (N)-PDU; and the other (N)-PDU and (N-1)-SDU). Thus accompanying mechanisms are required to keep track of associated security labels during this break-down of (N)-SDU both in the case of variable-(set-)labelling (when a single label or set of labels is associated with an (N)-SDU) and multi-(set-)-labelling (when labels or sets of labels are associated with different parts of an (N)-SDU)).

The structuring of SDUs (into one or more labelling fields) is not currently recognized in any of the relevant layer service definitions and thus, the chances are small that ready-made mechanism implementation will exist that keeps track of multi-(set-)labelling. Similarly the labelling of SDUs is not recognized and so ready-made mechanism implementations are unlikely to exist which support variable-(set-)labelling.

The implementation of these mechanisms must have its function verified since it is relied upon to maintain the separation of data associated with different security labels. To a certain extent the complexity of these implementations can be reduced by eliminating segmentation & reassembly, blocking & deblocking, and concatenation & separation mechanisms in each layer protocol. This would have the desirable effect of reducing the verification effort required, but may have the undesirable effect of reducing the functionality or efficiency that the protocols could provide.

If a variable- or multi- labelling (N)-facility uses a variable- or multi- set-labelling (N-1)-facility the set of labels used by the (N-1) facilities consists of only one element – the one associated with the (N)-facility.

It is not possible, in general, to support variable- or multi- set-labelling (N)-facility using a variable- or multi- labelling (N-1)-facility.

7. Summary of (N)-Facility Support

The following table shows the mechanisms proposed for the support of a labelling (N)-facility of one type by a labelling (N-1) facility of another.

			X-labelling (N-1)-facility					
	X		*F*	*V*	*M*	*Fs*	*Vs*	*Ms*
	fixed-	*F*	1	2	3	0+1	0+2	0+3
	variable-	*V*	5	6	7	0+4,0+5	0+6	0+7
X-labelling	multi-	*M*	5	8	9	0+4,0+5	0+8	0+9
(N)-facility	fixed-set-	*Fs*	–	–	–	1	2	3
	variable-set-	*Vs*	–	–	–	4,5	6	7
	multi-set-	*Ms*	–	–	–	4,5	8	9

Key to mechanism numbers:

–	there is no mechanism to provide this support
0	(N) label is (N-1) label set's only member
1	1:1 correspondence between labels for (N)- and (N-1)- connections
2	(N-1) SDU labels are fixed (N) label or label set
3	(N-1) labelled fields labels are fixed (N) label or label set
4	(N-1) label set encompasses all (N) labels used in a connection
5	(N)-connection split over (N-1)-connections according to (N) label, and/or (N)-connection separated into (N-1)-connections according to (N) label
6	labelled (N)-SDU is mapped to labelled (N-1)-SDU
7	labelled (N)-SDU is mapped to (N-1)-labelled field
8	(N)-labelled field is mapped to labelled (N-1)-SDU
9	(N)-labelled field is mapped to (N-1)-labelled field

There are two independent qualities of the (N)-facilities that can be isolated for consideration as "N" varies:

(1) uniform versus diverse labelled-field support

(fixed-(set-)labelling provides uniform labelled-field support, and variable- and multi- (set-)labelling provide diverse support);

(2) single label versus label set support

(fixed-, variable- and multi- labelling provide single label support, and fixed-, variable- and multi- set-labelling provide label set support).

Each type of support can continue to be supported if that type is provided by the layer below. Uniform labelled-fields can be supported on diverse ones and vice versa. Single labels can be supported on label sets but label sets cannot be supported on single labels. Thus, in principle, a different type of support for labelled-fields (uniform or diverse) could be provided at every layer – whereas, once single label support are provided, all the higher layers must also have single label support.

The choices for the provision of the required types of association can be expressed in terms of the position at which the transition is made from the type ((1) and (2) above) of the support supplied at the bottom layer to the type required at the top layer.

Since (going up the protocol stack) a transition from single label to label set support cannot be reversed, and label sets can be used where single labels are required, there is no case for making that transition. As such, all (N)-facilities are best chosen to support label sets.

In individual instances it is necessary to choose the positions of the transition from diverse to uniform label support given the requirement for different kinds of association (which will all be of the type to support label sets).

8. Support of a Fixed Label-Set Association

Going up the protocol stack, diverse labelled-fields are provided at the network layer domain access sublayer, and in order to provide uniform labelled-field support a transition must be made in some higher layer.

The (L+1)-layer is the lowest such layer. The benefit of providing the mechanism for supporting a fixed-set-labelling facility at as low a layer as possible is that local labelling operations on a per-SDU basis will not be required above it, only on a per-connection basis. In many computers this gives a realistic opportunity for providing a separate process per connection, and therefore of using process separation as the basis for verification of confinement.

The mechanism in the layers above the (L+1)-layer is as described above for the support of a fixed-set-labelling (N)-facility by a fixed-set-labelling (N-1) facility. An application association can then similarly be provided from a presentation layer connection.

9. Support of Variable or Multiple Label-Set Associations

For a multiple label-set association the application protocol in use must provide a multi-set-labelling application-facility. This can be provided by one or more fixed-variable- or multi-(set-)labelling (N)-facilities in layers below. As already mentioned, given the basic support for label sets there is little justification for using single labelling (N)-facilities. Support by each of these kinds of protocol are discussed in turn.

9.1 Support by variable- and multi- set-labelling facilities

As noted above when discussing the provision of variable- and multi- set-labelling (N)-facilities using variable- and multi- set-labelling (N-1)-facilities the mechanisms involved are liable to result in relatively complex implementations which also require verification. Therefore it is preferable to convert from diverse labelled-field to single labelled field as high up the protocol stack as possible to reduce the number of such steps required. Given that such conversion is possible at the very top of the stack there seems no good reason to supply variable- or multi- set-labelling facilities below the application layer.

9.2 Support by single-labelling facilities

Four categories of mechanism were discussed above in the provision of variable- and multi-set-labelling (N)-facilities using variable- and multi- set-labelling (N-1)-facilities:

- set labelling;
- splitting;
- connection-separation; and
- splitting and connection-separation.

The part of the latter three mechanisms used to manage the creation and deletion of the various connections; to synchronize them and to represent control protocol elements correctly, require a certain amount of complexity which may be duplicated in other layers. The benefit received is the potentially lower verification costs of lower layers (since confinement could be assured using process separation).

Splitting is a common function of the transport layer. Hence at this level one method of support by fixed-set-labelling protocols could map a variable- or multi- set-labelling transport protocol onto a fixed-set-labelling network protocol. However, performing this mapping at the transport layer gives very little benefit since the transport, session and presentation layer protocols will still require the complexity (and probably non-standard implementation) of multi-labelling protocols.

A similar argument holds at the session layer where connection-separation is a function of the layer.

Although splitting and connection-separation are not so common in the presentation layer a special purpose implementation could be provided which performed them. Furthermore the presentation layer contains a mechanism for distinguishing parts of the data associated with different labels – insofar as the presentation context could be used for this purpose. In order to use this method of label association an application service element must interpret application layer label information and select an appropriate presentation context.

However, with little additional effort an application service element could select different presentation connections (i.e. perform splitting) itself. It could also provide connection-separation. This has the additional benefit of allowing the use of standard implementations of both presentation and session protocols. If splitting and/or connection-separation are to be used, positioning it the application layer would give the greatest benefit.

Because of its nature, the use of connection-separation destroys any assurance of continuity of connection (one aspect of the maintenance of service security service) that might be provided by any lower level protocol (such a service can be provided at the transport layer). When connection-separation is used above the transport layer it also effectively prevents access to the frozen references function that could otherwise be made available.

If splitting is used at layer (N) without connection-separation a large number of simultaneous (N-1)-connections may be required to support a wide range of labels. This will have efficiency or economy disadvantages.

The complexity of splitting and/or connection-separation (by their nature trusted functions) means that the simpler alternative mechanism, set labelling, has much to recommend it. The association of fixed sets (perhaps ranges) of labels with each connection has the properties:

- the transfer of diversely labelled ("multi-class") data is possible;
- other than connection establishment very few existing implementations need supporting code changes;
- it is possible to use processes with a fixed associated set of labels to support (N)-connections (and thereby use process protection as part of the assurance of label separation);

- even when the label associated with data is known precisely it may be represented imprecisely in lower layers using a set of labels which include it;
- the code implementing such a process must be verified to ensure that it separates the labels in the fixed set.

The latter point would be a greater disadvantage if it were not for the likelihood of such code requiring verification for other purposes (such as maintenance of data integrity).

10. Implementation Choices and Rationale

The means of associating security labels with communicated data are, by and large, associated with matters that are purely local to a single open system.

If the open systems in question provides a TCB sufficiently flexible to maintain the separation of data associated with different security labels, even when they are manipulated by the same protocol entity, then it is preferable to use the TCB's facilities for labelling since this will require no special purpose protocol elements, no special purpose mechanisms between the application layer and the (L+1)-layer, and a lesser degree of implementation verification. In addition, this approach would allow the provision of a variable or multiple label(-set) association.

If protocol entities have to maintain the separation of data themselves then there is a benefit in reducing the number of connection instances which are required to deal with data associated with multiple security labels. This will reduce the extent of assurance required and improve the feasibility of formal verification.

To this end a requirement for fixed label associations should be met by a series of fixed-set-labelling facilities in diminishing layers with the final mapping to variable-set-labelling carried out by the user of the (L)-layer.

The requirement for variable or multiple label(-set) associations should be met by a variable- or multi- (set-)labelling application-facility mapping onto a fixed-set-labelling presentation-facility (and then proceeding as for the recommendation for fixed label associations).

Non-TCB reliant mechanism implementations that deal with multiple labels must represent data to security label (or label set) bindings explicitly. This requirement will mean that standard "off the shelf" protocol implementations are unlikely to be available for these mechanisms.

Mechanism implementations (e.g. processes based) that provide fixed-labelling (N)-facilities for a connection associated with a particular label, and which are not re-used for other labels, could make use of process-based data separation that a TCB may provide. This would result in the use of "off the shelf" protocols becoming feasible.

References

[ISO84] International Organization for Standardization (ISO), "Information processing systems – Open Systems Interconnection – basic reference model", ISO 7498, 15 October 1984

"Information processing systems – Open Systems Interconnection – basic reference model Technical corrigendum 1", ISO 7498/Cor.1, December 1988

"Security Labels in Open Systems" (Position Paper), Russell Housley (Xerox Corporation)

Security Labels in Open Systems: A Position Paper

Russell Housley
Xerox Corporation
McLean, Virginia

10 April 1990

Security Labels

In open systems, security labels tell the protocol processing how to handle the data. Security labels contain security attributes of data. Security attributes are those that state what protections that must be afforded the data, and they state how much confidence should be placed in the data.

Data confidence was originally called "integrity" by Biba[1]. The term confidence is used in this paper so that "Biba integrity" is not confused with the integrity security service described in the Organization of International Standardization's (ISO) Open Systems Interconnection (OSI) Security Architecture[2,3].

Traditionally, security labels have been used to state the sensitivity of the data. The protocol processing uses the sensitivity label to provide confidentiality. That is, to protect the data from unauthorized disclosure. For example, the transport protocol may choose to encrypt a connection in order to protect the data from disclosure.

Security labels may also be used to state the integrity of the data. The protocol processing uses the integrity label to provide integrity. That is, to protect the data from unauthorized modification. For example, transport protocols may choose between two error detection code algorithms based on the integrity label.

Security labels may also be used to state the confidence that should be placed in the data. Confidence labels are fundamentally different than sensitivity and integrity labels; they are not associated with any of the security services described in the OSI Security Architecture. The protocol processing should preserve the data confidence. For example, routers may choose a particular path through the network to preserve the data confidence.

Security labels may be used to make rule-based access control (RBAC) decisions. Sensitivity labels, integrity labels, and/or confidence labels may each be involved in the access control decision depending on the security policy being enforced.

Other Labels

Recent labeling discussions have included availability labels[4], authorization codes, and billing codes[5].

Availability labels denote the accessibility of the data. For example, payroll data must be available with sufficient lead time to print the checks. Availability, although important, is not an attribute which belongs in the security label. Availability requirements are currently met through the use of quality of service (QOS) and precedence parameters.

Authorization codes tell whether or not a particular user is permitted to use network resources. Again, authorization codes are important, but they should not be included in security labels. Authorization codes describe the permissions granted to a particular user or group of users; they do not tell the protocol processing how to handle the data.

Billing codes tell who should be billed for the network resources which are consumed. Like authorization codes, billing codes do not tell the protocol processing how to handle the data, so they should not be included in security labels.

End System Security Label Requirements

Some operating systems label the data they process. Some database management systems (DBMSs) perform similar labeling. The format of these labels is a local matter.

Trusted systems which implement RBAC policies require labels on the data they import. The labels permit the trusted system to perform trusted demultiplexing. That is, the network traffic is given to a process only if it has sufficient authorization for the data. In many cases, the trusted system must first translate the network security label into the local form before it can make the access control decision.

When two end systems communicate across a network, common label syntax and semantics are needed. The label must communicate all of the data handling requirements between the two communicating end systems.

Intermediate System Security Label Requirements

Intermediate systems, commonly called routers, make routing choices or discard traffic based on the security label. Bridges, packet switches, and application gateways also share this characteristic. For simplicity, the discussion will be limited to routers, but the concepts also apply to bridges, packet switches, and application gateways.

The security label used by the router should contain only enough information for the router to make its routing/discard decision. The label used by the router may very well be a subset of the security label used by the application. For example, copyright is not likely to affect routing.

References

[1] Biba, K. J. "Integrity Considerations for Secure Systems," ESD-TR-76-372 and MTR-3252. The MITRE Corporation, Bedford, MA, April 1977.

[2] ISO 7498, Information Processing Systems - Open Systems Interconnection - Basic Reference Model.

[3] ISO 7498/2, Addendum to ISO 7498 on Security Architecture.

[4] Branstad, Dennis K. "Categories of Information Requiring Protection," Proceedings from INTEROP 89, Advanced Computing Environments, October 1989.

[5] Estrin,Deborah. "Requirements for a Commercial-Use IP Security Option" Proceedings from INTEROP 89, Advanced Computing Environments, October 1989.

"Security Labels in Open Systems" (Presentation Slides), Russell Housley (Xerox Corporation)

Security Labels in Open Systems

Russell Housley
Xerox Special Information Systems
McLean, Virginia

31 May 1990

Data Security

- The measures taken to protect data from accidental, unauthorized, intentional, or malicious modification, destruction, or disclosure.

- A condition that results from the establishment and maintenance of protective measures.

Xerox Special Information Systems

Security Labels

- Security labels tell the protocol processing how to handle data communicated between open systems. That is, the security label indicates what measures need to be taken to preserve the condition of security.

- "Handle" denotes the activities performed on data such as collecting, processing, transferring, storing, retrieving, sorting, transmitting, disseminating, and controlling.

Protection from Modification, Destruction, and Disclosure

Protection from writing and deleting:

- Data integrity service
- Biba integrity

Protection from reading:

- Data confidentiality service
- Bell & LaPadula secrecy

Xerox Special Information Systems

Integrity Labels

- Support rule-based access control (RBAC) policies
- Tell the amount of confidence that may be placed in the data
- Tell which measures the data requires for protection from modification and destruction

- Data may be relabelled with lower integrity label as a result of being handled by an entity with a lower integrity label

Xerox Special Information Systems

Sensitivity Labels

- Support rule-based access control (RBAC) policies
- Tell the amount of damage that will result from disclosure of the data
- Tell which measures the data requires for protection from disclosure

- Data may be relabelled with a higher sensitivity label as a result of being handled by an entity with a higher sensitivity label

Xerox Special Information Systems

Availability

- Availability deals with low time delay to access network resources
- Availability, for some applications, may be important for security
- Availability, however, is not an element of data security (protection from modification, destruction, or disclosure)
- Availability requirements can be meet by appropriate use of the Quality Of Service (QOS) and Priority protocol fields

Xerox Special Information Systems

Authorization and Billing Codes

- Authorization codes deal with access control of network users to network resources
- Billing codes deal with payment for the use of network resources
- Authorization and billing codes deal with access control of network users to network resources
- Authorization and billing codes are not an element of data security (protection from modification, destruction, or disclosure)
- Authorization and billing codes need protocol fields, but the security label is an inappropriate field for them

Xerox Special Information Systems

Two Major Types of Systems in OSI

- End Systems (ES)
- Intermediate Systems (IS)
  - For this discussion, ISs will include routers, packet switches, and bridges

- ESs and ISs have different security labels requirements

Xerox Special Information Systems

End System Label Requirements

- Between two ESs, confidentiality and integrity requirements must be exchanged with data
- Multilevel ESs on multilevel networks require security labels in order to perform trusted demultiplexing
- ESs usually translate network security labels to a local format

Xerox Special Information Systems

Intermediate System Label Requirements

- Security labels include enough information to make routing/discard decisions
- Labels used by ISs may be a subset of the security label used by the user process/application layer
- Few ISs in a network actually make label-based routing/discard decisions, so security label parsing should not be imposed on all ISs
- ISs do not usually translate network security labels to a local format

ES and IS Label Requirements

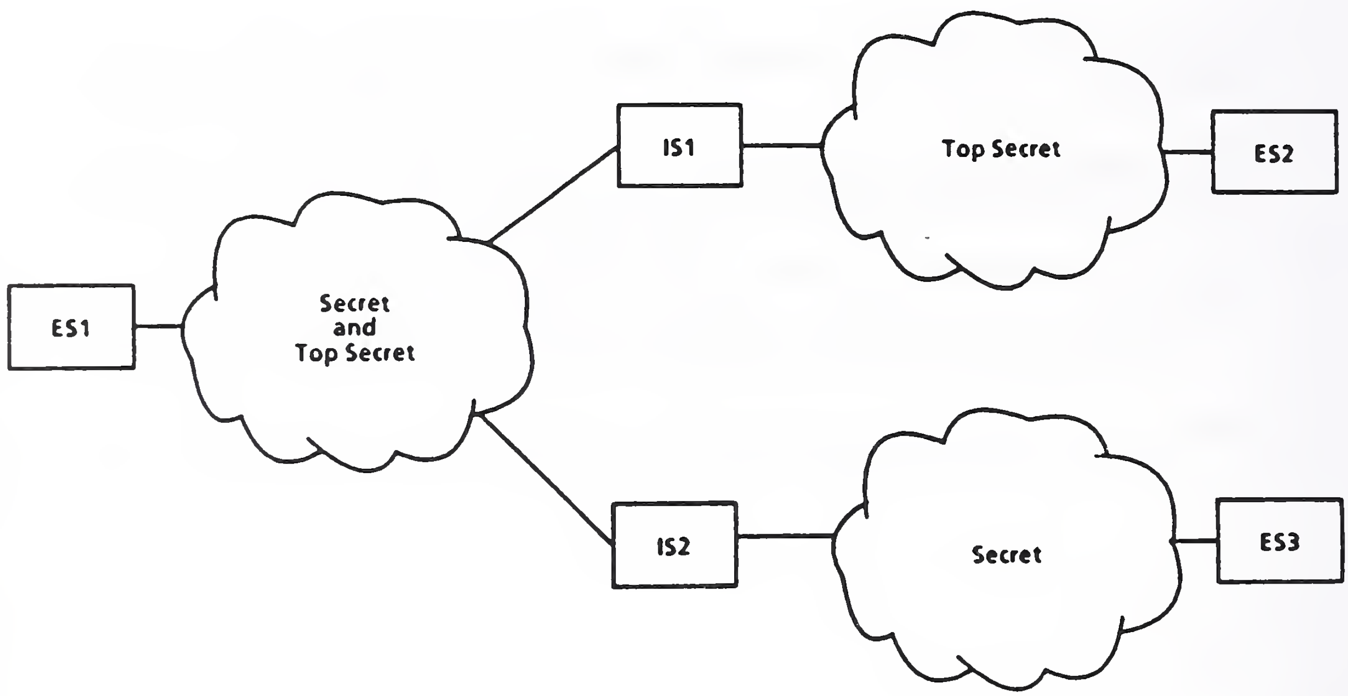

Xerox Special Information Systems

ES and IS Label Requirements

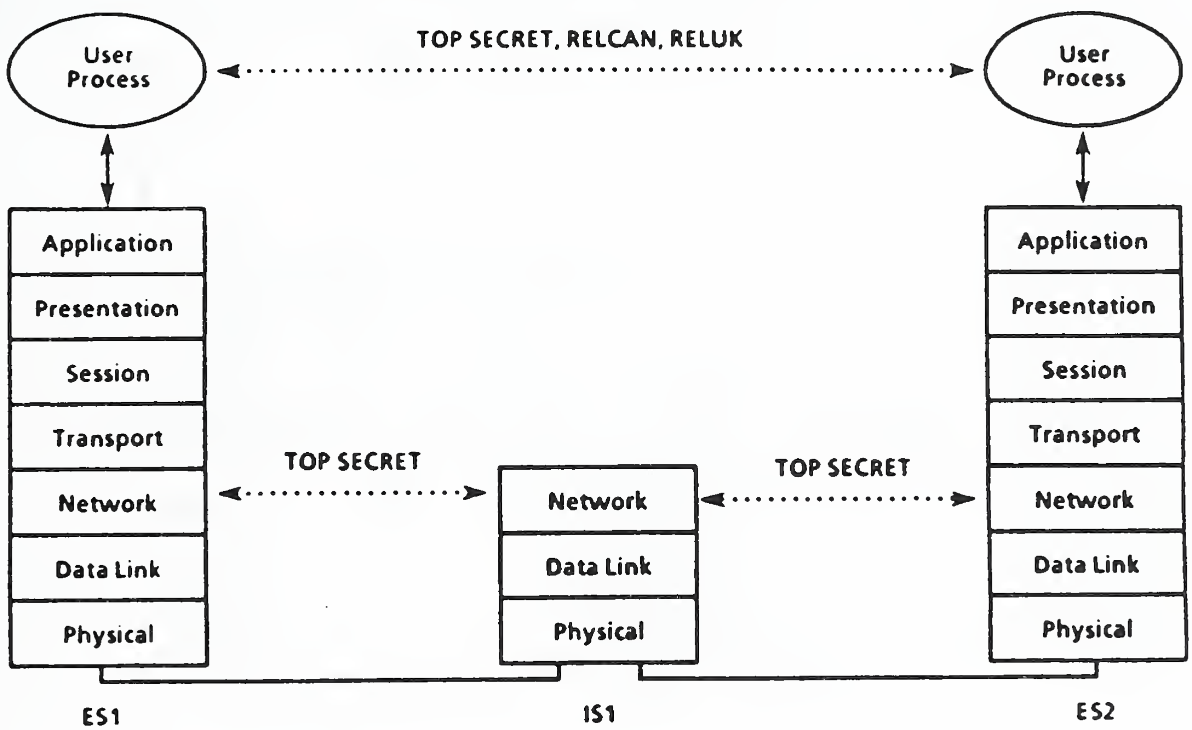

Xerox Special Information Systems

ES and IS Label Requirements

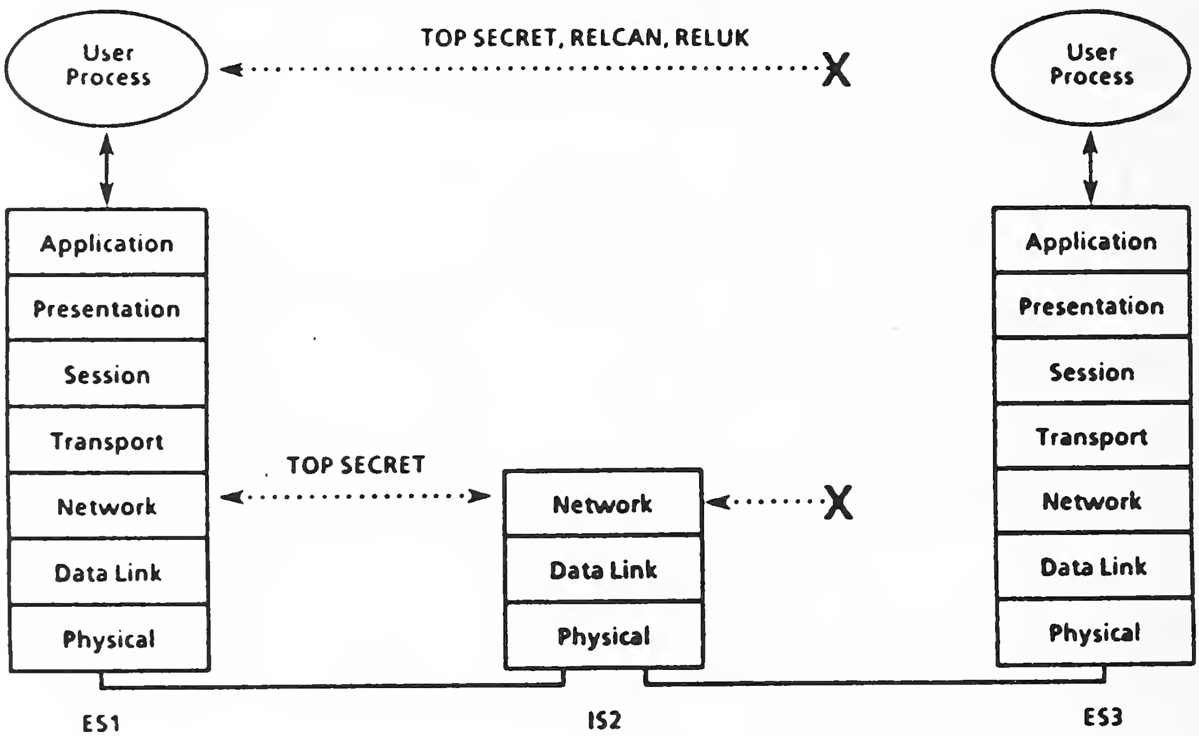

Xerox Special Information Systems

Approaches to Labelling

- Explicit vs. Implicit
- Connectionless vs. Connection-oriented

Xerox Special Information Systems

Explicit Labelling

- Bits in the Protocol Data Units (PDUs) give the label
  - Example: IP Security Option (IPSO)
- Can be used with both connectionless and connection-oriented labelling

Xerox Special Information Systems

Implicit Labelling

- Some attribute is used to determine the label
  - Example: Choice of SP4 cryptographic key

- Can be used with both connectionless and connection-oriented labelling

Connectionless Labelling

- Label every PDU
- Limited label size
- Limit may prohibit the label from meeting ES requirements
- Meets IS requirements

Xerox Special Information Systems

Connection-oriented Labelling

- Label virtual circuit/connection at establishment
- More compatible with ES requirements than IS requirements
  - May be compatible with X.25 Packet Switches
- Does not support connectionless protocols

Xerox Special Information Systems

Labelling Within the OSI Reference Model

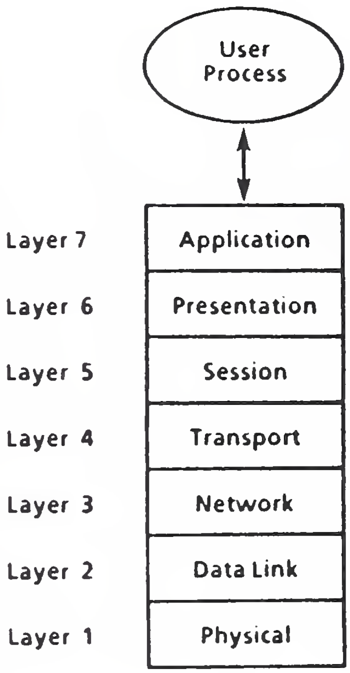

- Discuss security labels within each of the seven layers

- Start with layer 1 and work up

Xerox Special Information Systems

Physical Layer Labelling

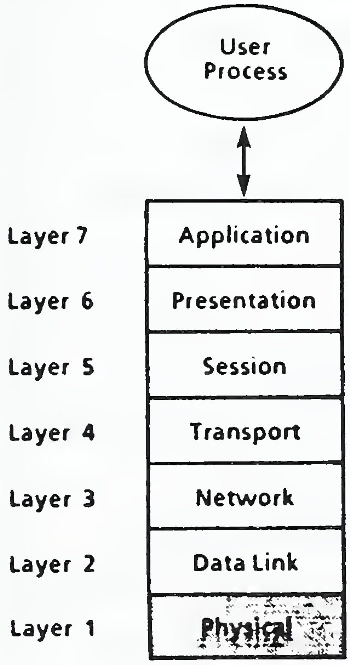

- Explicit labels not possible
  - No connectionless or connection-oriented labels

- Implicit labels possible

Xerox Special Information Systems

Data Link Layer Labelling

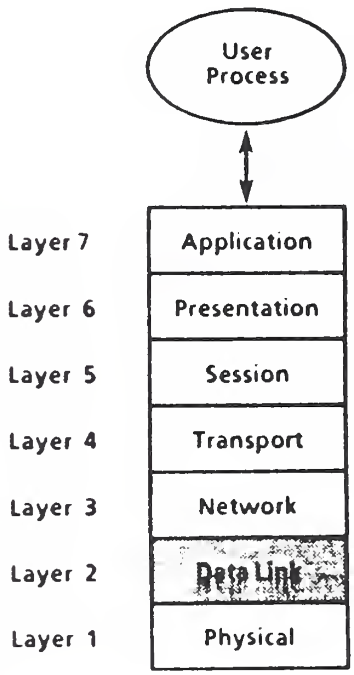

- Good for meeting IS label requirements
- Okay for meeting ES label requirements (with small labels)

- Explicit labels possible
  - Connectionless labels on each PDU possible
  - Connection-oriented labels possible for connection-oriented data link protocols (e.g., LLC Type II)
- Implicit labels possible

Xerox Special Information Systems

Network Layer Labelling

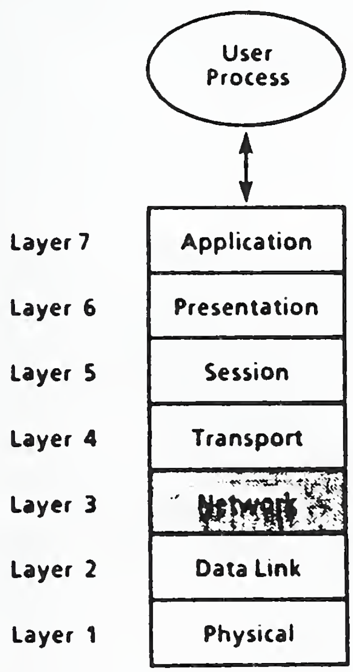

- Good for meeting IS label requirements
- Okay for meeting ES label requirements (with small labels)

- Explicit labels possible
  - Connectionless labels on each PDU possible
  - Connection-oriented labels possible for connection-oriented network protocols (e.g., X.25)
- Implicit labels possible

Xerox Special Information Systems

Transport Layer Labelling

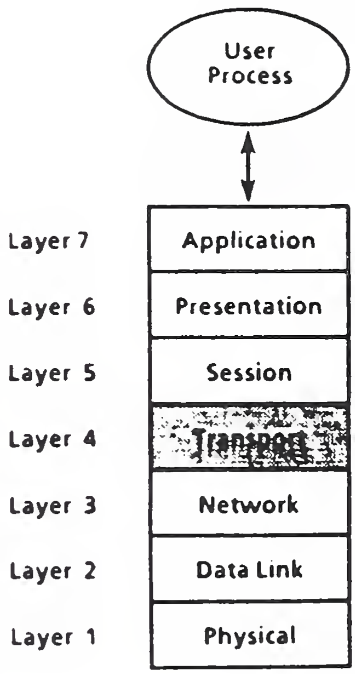

- Can not meet IS label requirements
- Good for meeting ES label requirements

- Explicit labels possible
  - Connectionless labels on each PDU possible
  - Connection-oriented labels possible for connection-oriented transport protocols (e.g., TP4)
- Implicit labels possible

Xerox Special Information Systems

Session Layer Labelling

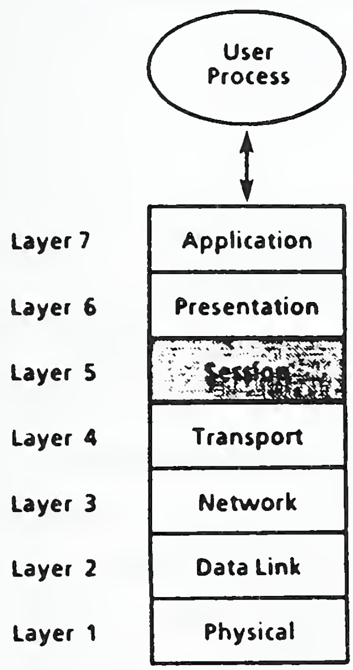

- Can not meet IS label requirements
- Poor for meeting ES label requirements (see IS 7498/2)
  - DNSIX *seems* to be doing session layer labels anyway

- Explicit labels possible
  - Connectionless labels on each PDU possible
  - Connection-oriented labels possible for connection-oriented session protocols (e.g., ISO Session)
- Implicit labels possible

Xerox Special Information Systems

Presentation Layer Labelling

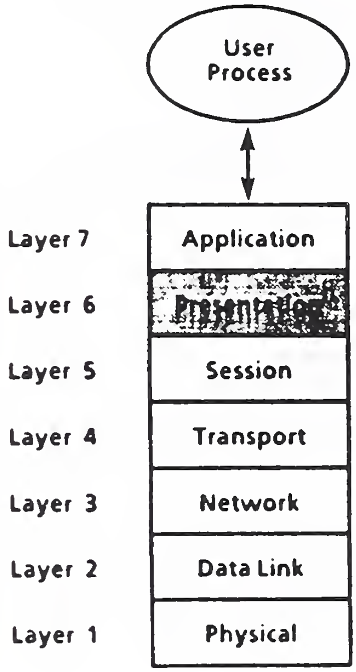

- Can not meet IS label requirements
- Good for meeting ES label requirements

- Explicit labels possible
  - Presentation syntax may include label
  - Naturally performs translation to local label format
  - connectionless or connection-oriented depending on the presentation protocol
- Implicit labels possible

Xerox Special Information Systems

Application Layer Labelling

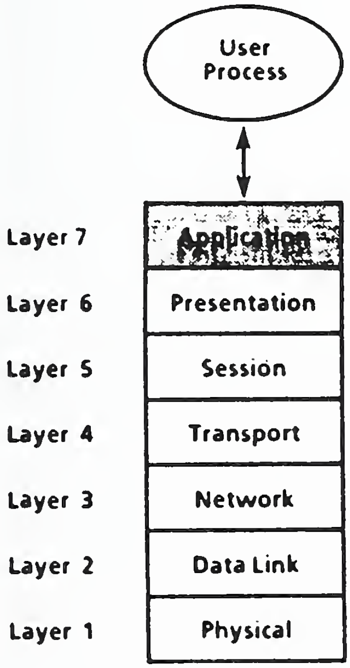

- Can not meet IS label requirements
- Good for meeting ES label requirements
- Explicit labels possible
  - Can include label information which is specific to a particular application without burdening other applications with syntax or semantics of that label
  - connectionless or connection-oriented depending on the presentation protocol
- Implicit labels possible

Xerox Special Information Systems

"Commercial IP Security Option" (Position Paper), John Linn (Digital Equipment Corporation)

COMMERCIAL IP SECURITY OPTION

John Linn
Secure Systems
Digital Equipment Corporation
85 Swanson Road, BXB1-2/D04
Boxborough, MA 01719-1326
Linn@ultra.enet.dec.com
Prepared for INTEROP '89 Invitational Workshop
October 1989, San Jose, California

1 Introduction, Context, and Assumptions

This note considers issues involved in definition of suitable IP security options to satisfy commercial market needs. The opinions expressed herein are those of the author, and do not represent official positions of Digital Equipment Corporation.

In comparison with the DoD environment, commercial environment definitions for clearances, data sensitivity labels, sensitivity categories, and data labeling are fragmented among larger numbers of organizations. The usage of, and supporting infrastructure for, rule-based access control (RBAC) is less established and mature in the commercial marketplace than in the DoD sector. In the commercial realm, clearances assigned by one organization are not generally transferable to, or interchangeable with, those of other organizations, although particular translations may be possible in the context of pairwise inter-organizational agreements. Similarly, sensitivity labels and category definitions are not generally interchangeable across organizational boundaries, although particular mappings may be possible given suitable agreements.

A few basic assumptions:

1. The CIPSO must accomodate definition of security policies and access attributes by customer organizations; insofar as feasible, policy definition should not be imposed or constrained by developers of equipment which generates or processes the CIPSO.
2. The CIPSO's role is to represent attributes used as inputs to RBAC decisions, identifying the sensitivity of a datagram's contents. In ECMA TR/46 ("Security in Open Systems: A Security Framework", July 1988) terminology, the CIPSO carries control attributes rather than privilege attributes.
   - Authorization mechanisms, establishing whether a host (or its users) are permitted to process information of a particular designated sensitivity, are outside the scope of the CIPSO.
   - Authentication mechanisms, serving to authenticate the identity of a particular host or user, are outside the scope of the CIPSO.
   - The CIPSO is not intended as a means to transport identity-based access control data structures (e.g., ACLs) to be associated with datagrams; interpretation of such data on a per-datagram basis is not generally appropriate.
3. Access decisions based on CIPSO-represented attributes can be made by several types of entities involved in communications processing, including:
   - reference monitors within communicating peer hosts

- security devices associated with communicating peer hosts
- intermediate systems such as gateways (feasible only if the CIPSO is readable as it traverses the intermediate system; inadvisable unless the CIPSO's integrity can be assured at the intermediate system)

4. While most datagrams will probably be labeled according to the conventions of a single labeling authority, the CIPSO should allow multiple authorities' labels to be distinguished and applied to a single datagram.
5. The integrity of the CIPSO's contents, and of its binding with a corresponding datagram, must be maintained from the point when the CIPSO is applied until the datagram is processed at its recipient system.

2 CIPSO Issues and Requirements

2.1 When Is a CIPSO Used?

It is a customer prerogative, outside the scope of the CIPSO standard, to dictate the circumstances under which CIPSO usage is required and which labeling authorities' CIPSOs should be applied to particular datagrams or associations.

While a CIPSO, labeling individual IP datagrams, is an important element in supporting rule-based security policies, it is not the only element and its use will not always be obligatory. The primary need for CIPSO per-datagram labeling arises in cases when an entity needing to make an RBAC authorization decision is unable to determine the access class of a datagram based on state information available to the determining entity; clear examples include:

- connectionless communications
- connection-oriented communications in which data of varying access classes may be carried on a single connection
- mediation of RBAC policies by intermediate systems (e.g., routers)

Many entities and channels will have fixed access class designations, allowing implicit labeling. In some other cases, implicit labeling can be achieved on a per-association basis, based on bindings established in the course of association initiation procedures.

Further, it is important to remember that customer requirements for security do not always imply requirements for rule-based security policies. Identity-based policies, at the granularity of hosts and/or users, can satisfy many customer needs in and of themselves, in a decentralized fashion without need for the centralized infrastructure needed to support RBAC. Many customer requirements will be best addressed with hybrid approaches employing both rule-based and identity-based policies.

2.2 Labeling authorities

In the commercial environment, labeling authorities correspond to customer organizations, organizational subunits, or established consortia thereof (e.g., the set of corporate participants in a joint venture). Labeling authorities define and coordinate the infrastructure (assignment of clearances to entities and of access class designations to data) which underlies RBAC. Related observations:

- The number of labeling authorities is large and unpredictable.

- Labeling authorities do not correspond to equipment vendors, except in the special case when a vendor is acting as a customer for its own use.
- Labeling authorities may be related hierarchically (as in the case of a department within an organization), but may also overlap without hierarchical implication (a joint venture between company A and company B is not superior to either A or B, and the labels specific to that venture may be significant within only a small part of each company).
- A given peer entity may be able to emit and process labels with formats defined by more than one authority.
- No association between network address or number and applicable labeling authority can necessarily be assumed: datagrams labeled in accordance with different authorities may coexist on the same network, and the scope of a single authority may span multiple networks.

A registry is needed to assign unique identification numbers to labeling authorities, so that labels generated in accordance with individual authorities' conventions can be interpreted unambiguously. A procedure like that used for Ethernet address assignment, yielding 48-bit labeling authority IDs, might be appropriate. If a group of customers, or a group of vendors addressing needs of a defined customer community, can establish a common labeling infrastructure and representation appropriate to the group, a labeling authority ID can be assigned on behalf of the group.

Interoperability across labeling authority boundaries may occur in two basic ways:

1. By translation between the conventions of one authority and the conventions of another, typically by relabeling based on pairwise inter-authority agreements. Such translation is likely to be performed at a relatively small number of translation points witin an authority's jurisdiction. Interorganizational agreements may constrain the set of policy translation points.
2. By labeling in accordance with the convention of a common authority recognized by both communicating peers: given domains D_A, D_B, and D_C, and peers P1 (a member of D_A and D_B) and P2 (a member of D_A and D_C) P1 may be able to communicate with P2 based on labeling information defined at the level of D_A. Operationally, this is a more convenient approach, even though it may not permit representation of the same granularity of attribute information as is achievable within the peers' unshared domains.

It's worth observing that the management burden attendant to supporting large-scale interdomain operations will be much less severe (avoiding n-squared problems of pairwise agreements between domains) where interdomain communications can be carried out under the rubric of a single shared domain (option 2 above), limiting the need for translation of attributes.

2.3 Access attributes and relations

The process of making an RBAC access decision on an individual datagram is a Boolean function of two types of inputs:

1. The sensitivity designation (access class) of the datagram's contents, as reflected in the CIPSO
2. The access rights of an entity

   •

A sufficiently rich "toolbox" of primitive access attribute operators should be available to satisfy the marketplace's anticipated needs, present and future. Different customer needs can suggest different types of access decision functions, broader than the set implied by the TCSEC concept of hierarchic levels and non-hierarchic categories as derived from the DoD clearance lattice and the Bell-LaPadula model. Possible relations include:

1. OR of access rights, in which the possession of any of a set of rights confers authorization to access particular data (e.g., access to funds transfer data by FINANCIAL or AUDITOR entities).
2. AND of access rights, in which all of a set of rights must be held in order to access particular data (e.g., access to payroll data only by entities holding FINANCIAL and PERSONAL rights).
3. EXCLUSION of particular entities from access based on certain of their RBAC attributes, even though access might be permissible based on other attributes. Binding of such attributes with an entity (a binding which the entity would be unable to revoke) would act to constrain the entity's access rights rather than expanding its privileges. An example usage might be restricting dissemination of US_EXPORT_CONTROLLED data to appropriate destinations based on absence of a NON_US attribute associated with an entity, instead of requiring that a US attribute be bound to all domestic host computers.
4. DOMINANCE, in which the right to receive information of a particular access class implies the right to receive information of dominated access classes.
5. CONFINEMENT, in which the right to emit information of a particular access class is restricted to entities whose access classes are dominated by the emitted access class.

Examples 4 and 5 illustrate the fact that RBAC decisions may not be symmetric; authorization to receive data with a particular CIPSO value doesn't necessarily confer authorization to emit data with the same value.

Some subtleties arise in encoding entities' access rights; for example, a host computer storing personal data (category P) as well as financial data (category F) might be capable of segregating those sensitivity categories, in which case it would be authorized to emit labeled datagrams carrying any of three types of designations:

- cat-P
- cat-F
- cat-P AND cat-F

Another host computer might be incapable of reliably separating the categories internally, in which case all of its emitted datagrams would be labeled:

- cat-P AND cat-F

2.4 Policy engine encoding concept

Variations among different domains' policies preclude comprehensive standardization of attribute definitions, yet it is important to help the cause of interoperability as far as possible. This subsection suggests an avenue for satisfying this goal, though detailed specification is a matter for further study.

For security components to satisfy the individual RBAC needs of different commercial customers, yet also provide the economies of scale associated with standard products, it is useful to develop components which can operate as customer-independent "policy engines". Such "engines" would perform attribute interpretation and access mediation based on several types of inputs:

1. Customer-provided encodings defining the space of access classes meaningful within that customer's scope and the relations among those classes
2. Per-entity encodings, defining the access classes within a customer's space which a given entity is authorized to emit and/or receive
3. Per-datagram CIPSO labels, defining the access class attributes of particular datagrams

Definition and use of an attribute encoding language would provide a level of indirection and abstraction. This could allow standard components to be tailored to different customers' needs by reconfiguration rather than reimplementation. Different reference monitors, implementing different policies, could be provided by loading of different customer-provided policy configurations. Language standardization would allow a customer to achieve interoperability across security components provided by multiple vendors.

3 Proposed Contents of CIPSO

I propose that only one CIPSO option be defined, in contrast to the pair of basic and extended options defined for the DoD environment in RFC-1038. Given:

- the absence of a uniform definition for security level across the commercial environment, and
- the incompatibility between the large anticipated number of labeling authorities and the bitmap protection authority flag representation used in RFC-1038

the Basic Option concept becomes vestigial. At best, a Commercial Basic option would provide a list of labeling authorities.

One could argue that inclusion of a labeling authority list in a basic option would allow an entity to reject a datagram without having to process an extended option containing contents defined by a labeling authority which the entity did not recognize. I don't believe, however, that the slight added effort involved in extracting the labeling authority IDs from well-known positions within authority-specific options warrants the redundant inclusion of an authority list within a Commercial Basic Option.

I propose that a Commercial Security Option, inspired by the DoD Extended Security Option, should contain the following elements:

1. Type code
2. Length indicator
3. Labeling authority ID
4. Additional security information, as defined by the labeling authority. Representation of this information in a standard encoding is recommended but not mandated.

As with the DoD Extended Security Option, the Commercial Security Option must be copied on fragmentation and may appear multiple times within a datagram, corresponding to labels applied in accordance with multiple authorities.

"Commercial IP Security Option" (Presentation Paper), John Linn (Digital Equipment Corporation)

Commercial IP Security Option: Issues and Concepts

John Linn

Secure Systems Development

Digital Equipment Corporation

Boxborough, MA

30 May 1990

Why a CIPSO?

- Administratively heterogeneous environment demands flexible labeling approach
- Labeling information appropriate for processing at ESs and ISs
- Even though commercial RBAC requirements are still emerging, a framework should be available to satisfy those requirements
- RFC-1038 accomodates only U.S. Government accrediting authorities

Policy and Scope Assumptions

- Accomodate policy definition by customer organizations; impose minimal constraints on policy characteristics
- CIPSO labeling's job: identify RBAC datagram attributes
  - -- not authorization, authentication data
  - -- not IBAC data structures (e.g., ACLs)
- Multiple authorities' labels should be applicable to a single datagram
- Some entities will be authorized to map between labels of different authorities, based on inter-authority agreements

Labeling Authorities

- Number of authorities is large and unpredictable
- Generally, authorities correspond to customers (or consortia thereof) rather than vendors
- Can't assume hierarchic relationship between authorities
- Can't assume that labeling authority can be identified implicitly based on address of entity emitting a label

Access Attributes and Relations

- Per-datagram RBAC decisions are Boolean functions of two types of inputs:
  - -- sensitivity designation (CIPSO label) of datagram
  - -- entity access rights
- Suggested CIPSO goals:
  - -- customer-definable decision functions, not just sensitivity designations
  - -- support multiple customer policies without reprogramming components
- Possible approach: define, interpret attribute language

Example Primitive Operators

- OR: "need FINANCIAL or AUDITOR rights to access FUNDS_TRANSFER data"
- AND: "need FINANCIAL and PERSONAL rights to access PAYROLL information"
- EXCLUSION: "no US_EXPORT_CONTROLLED data to NON_US entities"
- RANGE: entity processes data in hierarchic level range
- DOMINANCE/CONFINEMENT: authorization to process data of particular access class carries implications about other access classes

Levels of Extensibility

- RFC-1038
- Bell-LaPadula with alternative authorities
  - -- would allow different domains, analogous policies
- "Policy engine"
  - -- vendor-implemented encoding with toolbox primitives
  - -- customer-defined attribute space, entity assignments
- Arbitrary customer policy definition
  - -- would require per-customer programming and limit mechanized interoperability

Candidate Option Format

- Four elements:
  - -- type code
  - -- length indicator
  - -- labeling authority ID
  - -- security information, as defined by labeling authority
- Must be copied on fragmentation
- May appear multiple times in a datagram, corresponding to different authorities' labels

"Security Labels Position Paper", Bill Maimone (Oracle Corporation)

Security Labels Position Paper

Bill Malmone
Oracle Corporation

Introduction

This position paper outlines the role of secrecy labels as used by a hierarchically subsetted database management system and the subsequent requirements for standardization of labels.

The hierarchical subsetting approach, based on early research on extensible trusted computing bases (TCBs), allows the DBMS to rely entirely on the host processing environment for enforcement of mandatory access control (MAC). A corollary of this approach is that the resulting DBMS product is portable to a variety of secure hosts, and that the DBMS inherits the same features or flaws related to the support of labels present in the host environment. These factors combine to magnify the importance of label standardization across hetergeneous secure operating systems.

The Role of Labels

While it is possible for a hierarchically subsetted database to rely completely on the mandatory trusted computing base (m-TCB) for most issues relating to mandatory security, the unique role played by labels in a secure environment requires special attention.

Labels are a fundamental requirement for enforcing a mandatory security policy, but their central role in policy enforcement dictates non-policy uses in non-mandatory enforcing TCB subsets. The presence of labels on the operating system storage objects used to store records in a multi-level secure (MLS) database naturally leads to queries (where authorized) like "show me the classification of all tuples returned", or "show me all tuples classified above secret". The properties of the security policy that the labels are used to enforce, in which the label on a particular datum determines whether it can be read or written, also makes a label useful in supporting a coherent user-interface. Database records might be highlighted on a screen according to label, or a user might be prevented (with a meaningful notification) from attempting to perform an update that would only be rejected by the MAC enforcement in the operating system. In each of these examples, the mandatory security label provided to the user by the DBMS is termed an advisory label, since it plays no actual policy enforcement role once it has passed outside the m-TCB.

Labels can clearly play a useful role in a database outside the m-TCB, but the full utility of a label cannot be realized simply by passing on a label retrieved from the m-TCB. Labels invariably have two forms: an internal, typically binary, format used for efficiency; and an external, human readable, format for users. While the need to connect heterogeneous secure operating systems will eventually prompt some level of standardization on internal formats, the DBMS is motivated to find standards in both internal and external forms. This is especially true for DBMS products which are portable across many trusted platforms and/or support distributed MLS databases.

External Format

The DBMS must provide standard human-readable labels in order to support a portable SQL interface. Without a standard form and a mechanism for maintaining standard meanings, SQL statements and applications using labels will not even be portable across homogeneous secure operating systems.

A global label naming service should be supported in MLS environments. If category number fifty is printed as "ZETA" on one site, than category fifty (or whatever fifty is translated to by a secure network) must be printed as "ZETA" on any other sites. A failure to do so will inhibit the portability of applications referencing such identifiers, or the ability to consolidate data from distributed databases.

An alias or short form should be a part of any external label standard. Display size affects a DBMS more acutely than an operating system, since tuple labels will often be displayed aside the corresponding tuple. A label that takes up most of a screen leaves little room to display the tuple being labelled.

Internal Format

The format of an internal label is somewhat less important to a DBMS outside the m-TCB than external formats, but again standards would be productive. Internal format differences are slightly less debilitating as they could be handled via port-specific datatypes and modules. While inconvenient to the vendor forced to implement the port-specific module, the number of DBMS porters is at least less than the total number of SQL users. For DBMSs which store labels internal to the database lack of standards makes portability and migration of data/databases extremely difficult.

Label Operations

Similarly, the types of operations supported by MLS operating systems (when requried by a DBMS or other application) should also be standardized to provide for better application portability, heterogeneous environment support, and more consistent implementation and user interfaces.

The following operations should be standardized: return the internal label on a subject or storage object, format an internal label into an external label, compare internal labels for dominance, compare internal labels for sort order, compute least upper or greatest lower bound.

"Security Labels, End-to-End Encryption, and Internets", Phil Mellinger (MITRE Corporation)

Security Labels, End-to-End Encryption, and Internets

Phil Mellinger, Networking Center

MITRE

Briefing Overview

- A Historical Perspective of Security Labels
- Overview of Security Label Usage in End-to-End Encryption (E3) Systems
- Labeling Lessons Learned and Issues

MITRE

A Historical Perspective

- DoD message-switching system (AUTODIN) modernization required newer packet-switched messaging systems (formerly I-S/A AMPE and now DMS) to support packet-switched messages at multiple security levels.

- Packet-switch messaging systems (I-S/A AMPE and DMS) required E3 (formerly IPLI and now BLACKER) to cryptographically segregate multiple security levels of packet-switched messages over single-level Defense packet-switching networks (DDN).

- Packet labels would allow multiple-security level packet-switched messaging system to identify the proper cryptographic community for a message to the E3 device.

MITRE

A Historical Perspective (continued)

- Optimal location for packet labels was the DoD Internet Protocol layer and DCA-proposed format (RFC 1038) was deemed acceptable to BLACKER's designers (NSA).

- Thus, BLACKER and DoD hosts that interface to BLACKER are built to support these *Internet Protocol Security Option* (IPSO) labeling standards.

MITRE

DoD IP DATAGRAM TRANSITING PHASE I BFE

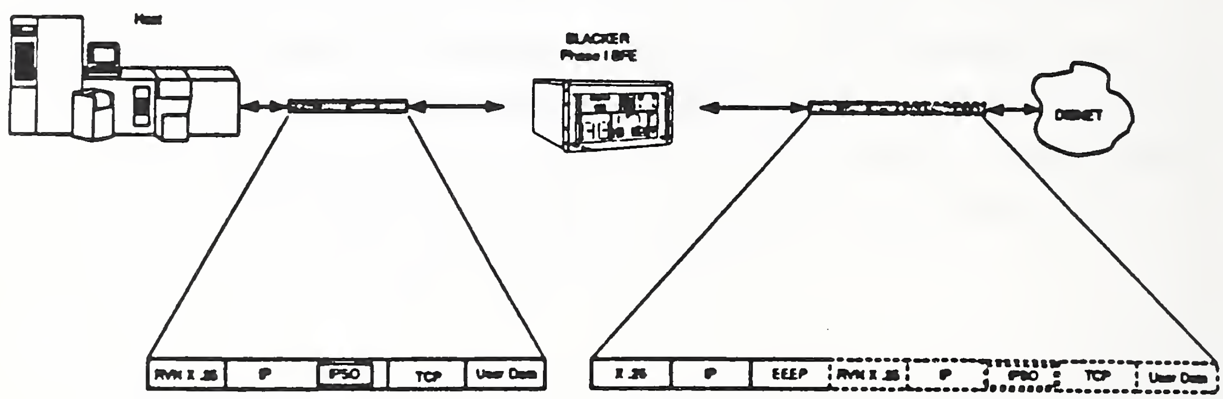

MITRE

Format of Basic Security Option

```
+------------+------------+------------+------------//----------+
:  10000010  :  XXXXXXXX  :  SSSSSSSS  : AAAAAAA[1]    AAAAAAA0 :
:            :            :            :        [0]             :
+------------+------------+------------+------------//----------+
 TYPE = 130     LENGTH    CLASSIFICATION       PROTECTION
                               LEVEL            AUTHORITY
                                                  FLAGS
```

MITRE

DDN-BLACKER Access Control

MAC Communities	GENSER	SIOP-ESI	SCI	NSA
Top Secret	Top Secret GENSER	Top Secret SIOP-ESI	Top Secret SCI	Top Secret NSA
Secret	Secret GENSER	Secret SIOP-ESI	Secret SCI	Secret NSA
Confidential	Confidential GENSER	Confidential SIOP-ESI	Confidential SCI	Confidential NSA
Unclassified	Unclassified GENSER	Unclassified SIOP-ESI	Unclassified SCI	Unclassified NSA

IF THE AIS PROCESSES	AND THE AIS'S CRITERIA CLASS IS	AND THE AIS'S SECURITY OPERATING MODE IS	THEN THE DAC GROUP IS	AND EMERGENCY MODE ENTRY IS
GENSER Information	Minimum of C2 (DDN Environment)	System High or Multilevel	Trusted/Open or DAA-dependent	DAA Discretion
		Dedicated	DAA-dependent	Prohibited
	Not Minimum of C2 (DDN Environment)	Dedicated	DAA-dependent	Prohibited

MITRE

DoD/GOSIP/X.25 UNITS TRANSITING PHASE III BFE

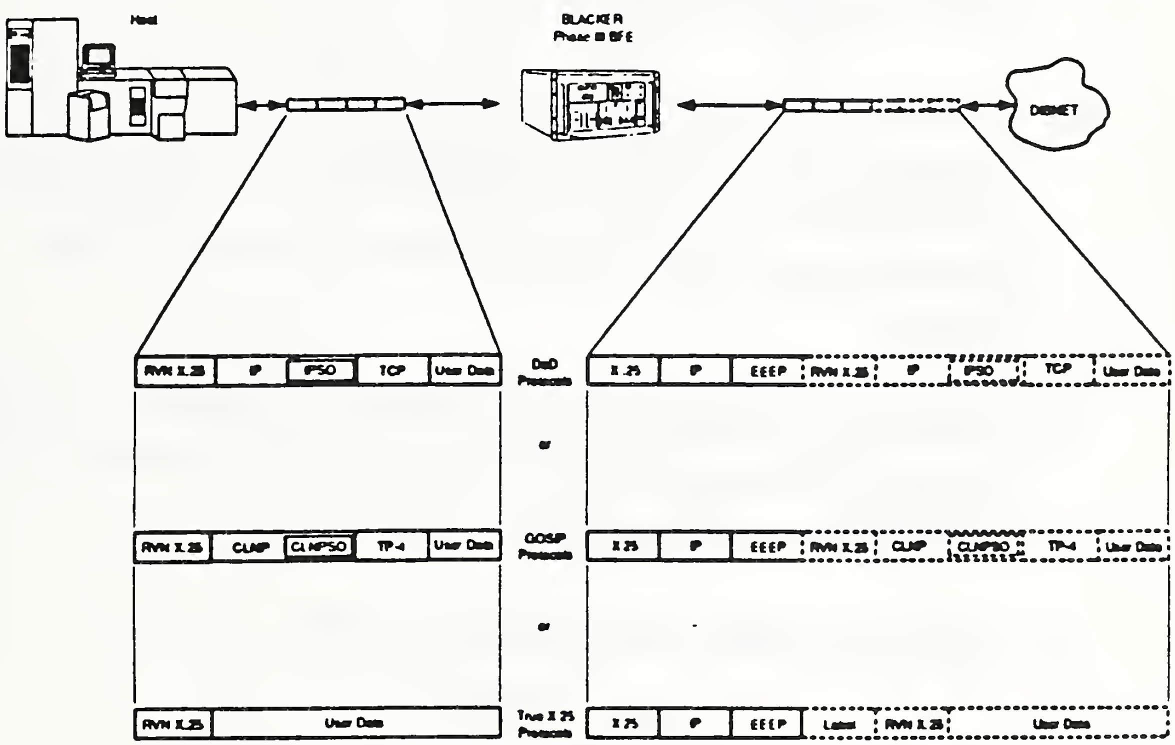

MITRE

Summary of BLACKER Labeling Functionality

- Labeling functionality used to identify security level of packet to BLACKER System.

- The security label contained in DoD IP (IPSO) and GOSIP CLNP (CLNPSO) is identical from parsing perspective to minimize code changes in BLACKER System (only label location differs between IP and CLNP).

- Labeling X.25 connections (Phase III) appears to require trusted X.25 machines, and was not instead left as an area for future research.

- X.25-only hosts must be single-level.

MITRE

Format of Extended Security Option

10000101	00001101	11111100	00000000...00000000
Option Type 133 (ESO)	Length of Option (13 Octets)	Authority Code (X-5000)	Information: All Zeroes (No Data)

Format Of The IP ESO Between Host And DoD-modified X-5000

10000101	00001110	11111100	00001010	10000010...00101011
Option Type 133 (ESO)	Option Length (14 Octets)	Authority Code (X-5000)	Length of Information (10 Octets)	Information: Key ID and IV (Data)

Format Of The CLNP ESO Between GOSIP-modified X-5000s

MITRE

Data Units Transiting Cipher X-5000-X.25/DDN

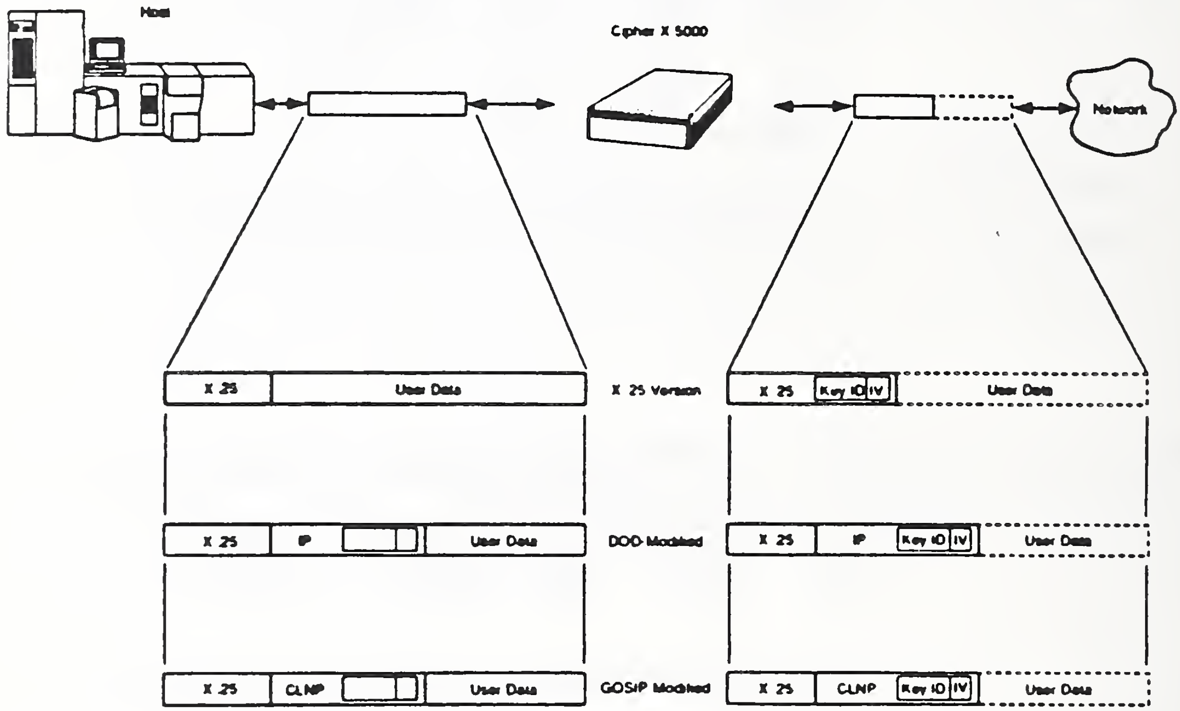

MITRE

Summary of Cipher X-5000 Labeling

- Each IP datagram's data field is encrypted (using a Key and a unique Message Indicator (MI)).
- Each IP datagram's header left unencrypted for routing through IP gateways.
- Encryption device places Key ID and MI in *host-provided* Internet Protocol Security Option (Type 133 - Extended) to avoid IP datagram expansion/fragmentation.

MITRE

Lessons Learned (Down-side)

- Large security labels may exceed maximum header size allowed (*minimize size of security labels*).

- Placement of security labels throughout 7-layer model could create tremendous overhead (standardize placement of security labels).

- Complex labels or multiple-labels are difficult to implement (*standardize format of security labels*).

- Flash-cutovers of large systems or internets to labeling are not pretty (*plan for transitioning from internets without labels to internets with labels*).

MITRE

Labeling Issues

- Who is in charge of security labels? (DCA, NSA, NIST, etc.)
- Why are security labels needed? (E3, COMPUSEC, etc.)
- Where should security labels be placed? (one layer or everywhere)
- How complex should security label handling be? (COMPUSEC good)
- What is the future -- trusted/COMPUSEC internets (with labels everywhere) or E3 communities over single-level (unlabeled) internets?

MITRE

"Position Paper", Nick Pope (Logica; U.K.)

NIST Workshop on Security Labels for Open Systems

Position Paper - Nick Pope (UK)

Work has been done in a number of areas in ISO on issues which relate to security labels for open systems. This includes work in the following areas:

- X.400 - Message Security Labels (X.411)
- Network and Transport Layer - Use of Security Labels for specifying Protection Quality if Service in an abstract form (SC6 WG2/WG4 Paris documents P2.38 and SC6/WG4/N581rev)
- Open System Security Frameworks - Security Domains (SC21 N4210)
- Open System Security Access Control Framework - Security Labelling (SC21 N4206)

Development of approaches to security labelling should take account of this work, and the results of any discussions at this workshop should be fed into ISO.

"Information Identification and Protection", Warren Schmitt (Sears Technology Services)

INFORMATION IDENTIFICATION

AND PROTECTION

WARREN SCHMITT
SEARS TECHNOLOGY SERVICES
IIP-1

ISSUE

- HOW DO WE MANAGE THE ASSET CALLED INFORMATION?
- RECOGNITION THAT SOME INFORMATION IS MORE VALUABLE THAT OTHER INFORMATION.
  - TRADE SECRETS
  - FINANCIAL TRANSACTIONS
  - ACQUISITION COSTS

INFORMATION IDENTIFICATION & PROTECTION

IIP-1A

ISSUE

- RECOGNITION THAT VALUABLE INFORMATION IS SUBJECT TO A VARIETY OF RISKS.
  - DESTRUCTION -- CUSTOMER FILES
  - MODIFICATION -- MGT. DECISIONS
  - DISCLOSURE -- COMPETITIVE ADVANTAGE

- CONTROLS NEED TO BE ESTABLISHED DURING APPLICATION DESIGN

INFORMATION IDENTIFICATION & PROTECTION

IIP-1B

DEFINITIONS

- RISK - VULNERABILITY - EXPOSURE
  - THE CONDITION OF BEING UNPROTECTED
- INFORMATION
  - KNOWLEDGE OR INTELIGENCE THAT IS REPRESENTED IN A COMPUTER BY DATA ,TEXT, VOICE, OR IMAGE.
- BUSINESS CONTROLS
  - COMBINATION OF ADMINISTRATIVE AND TECHNICAL CONTROLS THAT PROVIDE A REASONABLE ASSURANCE THAT THE ENTERPRISE'S OBJECTIVES WILL BE EXECUTED AS PLANNED.

INFORMATION IDENTIFICATION & PROTECTION

IIP-2

MAJOR RISKS

CONDITION	ACT
AVAILIBILITY	DESTRUCTION
INTEGRITY	MODIFICATION
CONFIDENTIALITY	DISCLOSURE

INFORMATION IDENTIFICATION & PROTECTION

IIP-3

CONCEPT

- A PROCESS WHERE BY WE CAN INDICATE THE INTRINSIC VALUE OF INFORMATION IN RELATION TO EACH OF THE THREE MAJOR RISKS

- AND BASED ON THE VALUE OF THE INFORMATION, IDENTIFY THE APPROPRIATE CONTROLS

INFORMATION IDENTIFICATION & PROTECTION

IIP-4

SIMILARITIES BETWEEN UNCLASSIFIED/COMMERCIAL

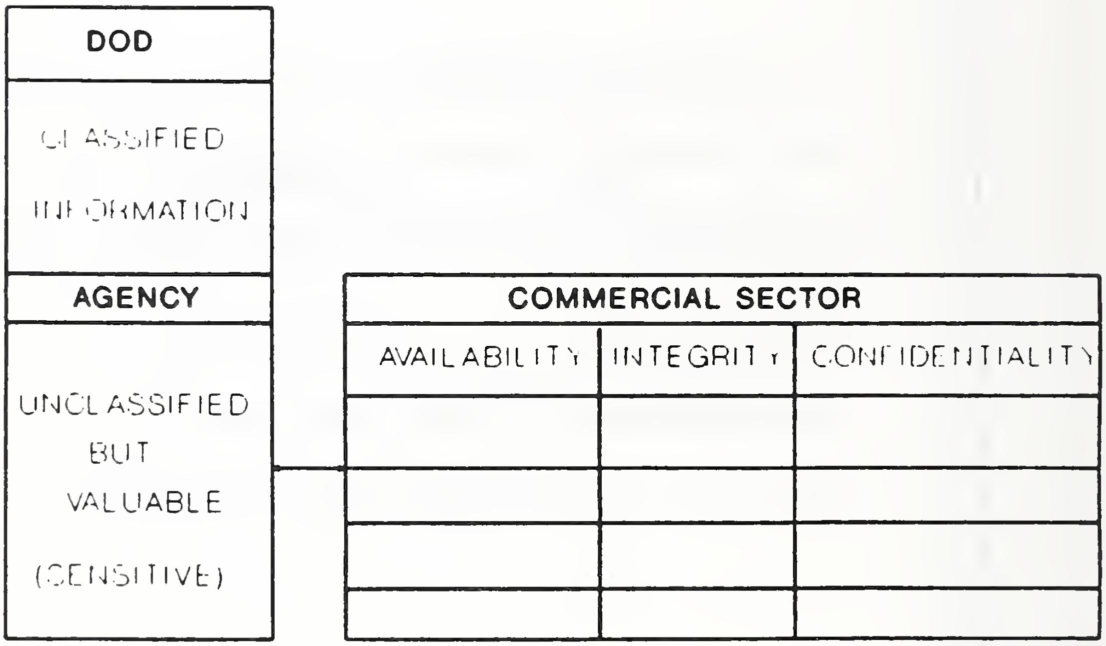

INFORMATION IDENTIFICATION & PROTECTION

IIP6

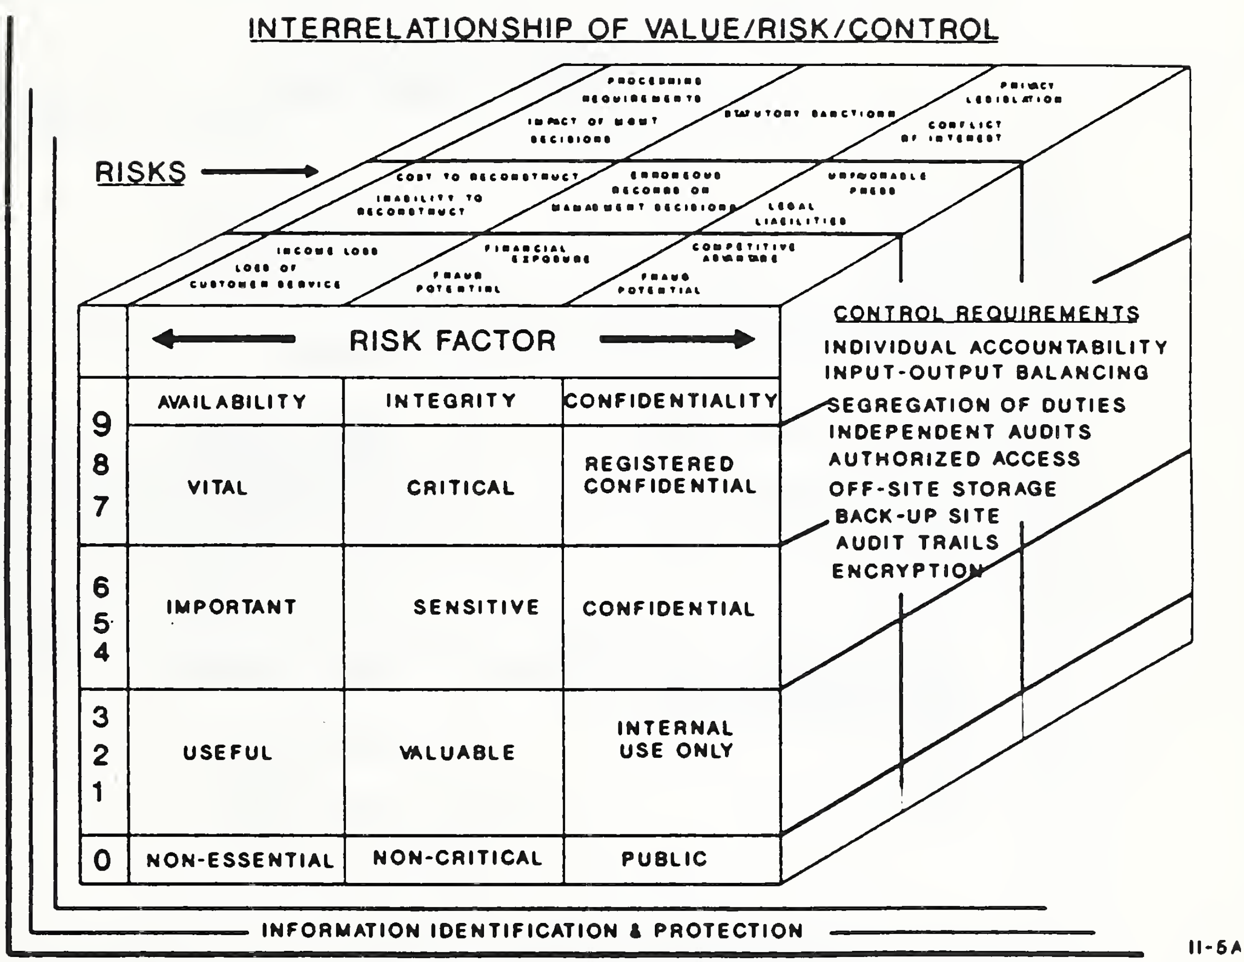
INTERRELATIONSHIP OF VALUE/RISK/CONTROL
RISKS
PROCESSING REQUIREMENTS
IMPACT OF MGMT DECISIONS
STATUTORY SANCTIONS
PRIVACY LEGISLATION
CONFLICT OF INTEREST
COST TO RECONSTRUCT
INABILITY TO RECONSTRUCT
ERRONEOUS RECORDS OR MANAGEMENT DECISIONS
UNFAVORABLE PRESS
LEGAL LIABILITIES
INCOME LOSS
LOSS OF CUSTOMER SERVICE
FINANCIAL EXPOSURE
FRAUD POTENTIAL
COMPETITIVE ADVANTAGE
FRAUD POTENTIAL
RISK FACTOR
AVAILABILITY
INTEGRITY
CONFIDENTIALITY
9 8 7
VITAL
CRITICAL
REGISTERED CONFIDENTIAL
6 5 4
IMPORTANT
SENSITIVE
CONFIDENTIAL
3 2 1
USEFUL
VALUABLE
INTERNAL USE ONLY
0
NON-ESSENTIAL
NON-CRITICAL
PUBLIC
CONTROL REQUIREMENTS
INDIVIDUAL ACCOUNTABILITY
INPUT-OUTPUT BALANCING
SEGREGATION OF DUTIES
INDEPENDENT AUDITS
AUTHORIZED ACCESS
OFF-SITE STORAGE
BACK-UP SITE
AUDIT TRAILS
ENCRYPTION
INFORMATION IDENTIFICATION & PROTECTION
II-5A

LABEL CONTENT

LABELS MUST, AT A MINIMUN, IDENTIFY EACH OF THE THREE VALUE/RISK RELATIONSHIP

AUTHORITY	DATE	VALUE / RISK RELATIONSHIP		
		A-5	I-7	C-9

AND WHEN APPROPRIATE, THE CONTROL REQUIREMENTS

INFORMATION IDENTIFICATION & PROTECTION

IIP-7

CLASSIFICATION CONTROL MATRIX
Risk/Exposure - Availability

FIGURE 1

TYPE OF RISK/EXPOSURE	INFORMATION	CLASSIFICATION	CONTROLS
HIGH	- STRATEGIC PLANS - INVESTMENT PORTFOLIO - CUSTOMER ORDER SYSTEM - AIRLINE RESERVATION SYSTEM - SWITCHED TELEPHONE NETWORK	VITAL	- ACTIVE HOT-SITE - OFF-SITE STORAGE - FIRE PROOF SAFE - HOURLY/DAILY UPDATES - TESTED CONTINGENCY PLAN (QRTLY) - INVENTORY OF OFF-SITE STORAGE (MONTHLY)
- LOSS OF INCOME AND CUSTOMER SERVICE (HOURS/DAYS) - INABILITY TO RECONSTRUCT - COST OF DIFFICULTY TO RECONSTRUCT - IMPACT ON MANAGEMENT DECISIONS - PROCESSING REQUIREMENTS HOURS/DAYS - CUSTOMER CONFIDENCE	- PRODUCTION ASSEMBLY LINE - PERSONNEL RECORDS - GENERAL LEDGER - CASH/DISBURSEMENT JOURNALS - CORPORATE DATA BASE - SOFTWARE CODE/DOCUMENTATION - SYSTEM GENS - OPERATING SYSTEMS - ACCESS SECURITY FILE - RESEARCH PROJECTS - LEGAL FILES/PROCEEDINGS	IMPORTANT	- IDENTIFIED ALTERNATE PROCESSING SITE - OFF-SITE STORAGE - FIRE PROOF SAFES - DAILY/WEEKLY UPDATES - TESTED CONTINGENCY PLAN (ANNUALLY) - RETENTION SCHEDULES - QRTLY. INVENTORY OF OFF-SITE STORAGE
	- MANUALS/PROCEDURES - DEPARTMENTAL/UNIT RECORDS (CURRENT) - RESEARCH MATERIAL (LIBRARIES) - HISTORICAL INFORMATION (1-3 YRS) - REGULATORY REQUIREMENTS	USEFUL	- LOCKED DESKS/CABINETS/ROOMS - DUPLICATE COPIES - PHYSICAL/ENVIRONMENTAL CONTROLS - RETENTION SCHEDULES
LOW	- HISTORICAL INFORMATION (3 YRS+)	NON-ESSENTIAL	- ORDINARY CARE

CLASSIFICATION CONTROL MATRIX
Risk/Exposure - Integrity

FIGURE 2

TYPE OF RISK/EXPOSURE	INFORMATION	CLASSIFICATION	CONTROLS
HIGH	- CORPORATE FINANCIAL RECORDS - INVESTMENT PORTFOLIO - SYSTEM GEN	CRITICAL	- ENCRYPTED TRANSMISSIONS - MESSAGE AUTHENTICATION - DUAL APPROVALS
- FRAUD POTENTIAL - FINANCIAL EXPOSURE - ERRONEOUS MGMT. DECISIONS - ERRONEOUS RECORDS - CUSTOMER CONFIDENCE	- SOFTWARE FOR CONTROLLED APPL. - ORDER RECORDS - SHIPPING RECORDS - PERSONNEL RECORDS - CASH RECEIPTS/DISBURSEMENTS - GENERAL LEDGER - LEGAL ACTIVITY - OPERATING SYSTEMS - CORPORATE DATA BASE - SOFTWARE DEVELOPMENT - OPERATING SYSTEMS	SENSITIVE	- INDIVIDUAL ACCOUNTABILITY - GUARDIAN - UPDATE AUTHORIZATION - AUDIT TRAILS - DATA ENTRY/OUTPUT BALANCING - SEGREGATION OF DUTIES - INDEPENDENT AUDIT - EDIT CHECKS - VERIFICATION/MONITORING - REASONABLENESS CHECKS - EXCEPTION CONTROL - PROGRAM LIBRARY CHANGE CONTROLS
	- DEPARTMENTAL/UNIT RECORDS - TIMEKEEPING RECORDS - DEPARTMENTAL CORRESPONDENCE - CORPORATE MANUALS - VENDOR MANUALS	VALUABLE	- GENERAL ACCESS/NEED-TO-KNOW - UPDATE/CHANGE PROCEDURES - RESPECT FOR COPYRIGHT REQUIREMENTS
LOW	- INTERNAL TELEPHONE DIRECTORIES - PUBLIC INFORMATION	NON-CRITICAL	- ORDINARY BUSINESS CONTROLS

CLASSIFICATION CONTROL MATRIX
Risk/Exposure - Confidentiality

FIGURE 3

TYPE OF RISK/EXPOSURE	INFORMATION	CLASSIFICATION	CONTROLS
HIGH	- GOVERNMENT CLASSIFIED INFORMATION - STRATEGIC PLANS - NEW DISTRIBUTION METHODS - FINANCIAL INFORMATION - EARNINGS FORECAST-LONGTERM - UNDERWRITING STRATEGIES - HIGH LEVEL CORPORATE REORG. - INVESTMENT STRATEGY - TRADE SECRETS - LEGAL ACTIVITY - EXECUTIVE PAYROLL/BONUS - PAYMENT	REGISTERED CONFIDENTIAL	- LABELING - CONTROL NUMBERS - RECORD OF EACH COPY - INDIVIDUAL ASSIGNMENT - HAND CARRIED - ENCRYPTION - ELECTRONIC - CONTROLLED DESTRUCTION - STORED IN SAFE - SENIOR MGMT. APPROVAL REQUIRED FOR ACCESS - ISOLATION OF COMPUTER SYSTEM

- COMPETITIVE VALUE
- FINANCIAL EXPOSURE
- FRAUD POTENTIAL
- LEGAL LIABILITY
- INSIDE TRADING
- UNFAVORABLE PRESS
- CONFLICT OF INTEREST
- VIOLATION OF PRIVACY
- CUSTOMER CONFIDENCE
	- CORPORATE DATA BASE
	- PROPRIETARY SOFTWARE
	- SHORT TERM MARKETING RESULTS
	- BID/PURCHASE ORDERS
	- REAL ESTATE SITE SELECTION
	- PRICING/RATE CHANGES
	- ORGANIZATIONAL CHARTS
	- ACCESS SECURITY PROFILES
	- CUSTOMER RECORDS
	- SOFTWARE DEVELOPMENT
	- PERSONNEL RECORDS
		CONFIDENTIAL	- INDIVIDUAL AUTHORIZATION
			- NEED-TO-KNOW/JOB RESPONSIBILITY
			- AUDIT TRAILS
			- NON-DISCLOSURE AGREEMENTS
			- PRIORITY MAIL DUTIES
			- NETWORK:
			  - LEASED LINES
			  - MESSAGE AUTHENTICATION
			- LOCKED DESKS/CABINETS
			- CONTROLLED DESTRUCTION

	- CORPORATE DIRECTIVES/CORRES. - MANUALS/PROCEDURES - TELEPHONE DIRECTORIES - TECHNICAL PAPERS	INTERNAL USE ONLY	- ACCESS/NEED-TO-KNOW - GENERAL ACCESS
LOW	- BOOKS, PERIODICALS - VENDOR MANUALS	PUBLIC	- ORDINARY BUSINESS CONTROLS

"Security Labels for the Defense Message System", Robert W. Shirey(MITRE Corporation)

SECURITY LABELS FOR THE DEFENSE MESSAGE SYSTEM

Robert W. Shirey

The MITRE Corporation, Mail Stop Z286
7525 Colshire Drive, McLean, Va. 22102-3481
703.883.7210, Shirey@MITRE.org

In 1988, the Department of Defense (DoD) began the Defense Message System (DMS) Program to improve and modernize (DoD) message handling systems. The DMS is defined to included "All hardware, software, procedures, standards, facilities, and personnel used to exchange messages electronically between organizations and individuals in the Department of Defense (DoD)." The DMS Target Architecture and Implementation Strategy (TAIS), a publicly-available document, defines how the system will evolve from the current baseline to the goal architecture.

The baseline includes (1) the Automatic Digital Network (AUTODIN) system, including local, baselevel elements; and (2) the electronic mail functions of the Department of Defense (DoD) internetworks, including Defense Data Network (DDN) long-haul backbones and their connected local area networks (LANs). The architectural goal, in brief, is to convert these systems to international X.400/X.500 standards and integrate them. This baseline is very large, and it is used by, and interoperates with, many other government agencies and the general Internet community. Thus, the DMS evolution will have effects outside DoD.

The Under Secretary for Acquisition has created a DMS Panel and a supporting DMS Implementation Group (IG), and has named the Defense Communications Agency as DMS Coordinator to Chair the IG. The IG prepared, and the Panel published, the TAIS, which covers three phases:

- Phase I Architecture for 1993: Telecommunication Automation
  - Reduce cost, automate functions, extend service to users
  - Phase in GOSIP protocols, X.400 messages, X.500 directory
  - Phase out existing protocols, procedures, and formats
  - Implement AUTODIN-to-DDN interfaces
- Phase II Architecture for 2000: SDNS Implementation
  - Consolidate existing message systems
  - Expand writer-to-reader connectivity, security, and support
  - Provide security based on Secure Data Network System standards
- Phase III Architecture for 2008: ISDN implementation

The Security Policy Working Group (SPWG) of the DMS IG is preparing high-level policies and plans to support the program. Among these will be a basic security policy that identifies minimum security safeguards that are required for operation of the DMS and for participation in it. The safeguards will implement a secure DMS with secure DMS components that perform the operational mission while minimizing the opportunity for sabotage, denial of service, data alteration or destruction, deliberate or inadvertent access to classified and sensitive unclassified information by unauthorized personnel, and unauthorized use of the system.

The DMS as a whole and DMS components that are networks are subject to the requirements of Enclosure 5 of DoD Directive 5200.28, Security Requirements for Automated Information Systems (AISs). For purposes of accreditation, Enclosure 5 ("Network Considerations") requires a network to be treated as either an interconnection of separately accredited AISs (which may be networks) or as a unified network. DMS components may be treated either way. However, the DMS as a whole will be treated as an interconnection of accredited AIS facilities, and the basic policy needs to specify interconnection rules.

To implement an interconnection policy, every DMS message will need to contain a standard security label when the message is transferred between components or is exported from the DMS. The DMS security architecture needs to specify the form and content of the label. The DMS security standard for component systems needs to specify how components use labels to determine how to handle messages. A DoD-wide labeling standard is needed that specifies uniform representation, syntax, and data structure for both human- and machine-readable forms of security labels in message handling (and in other communication protocols), and that specifies algorithms for the mappings between these forms.

The MITRE Corporation, sponsored by the Director of Information Systems, Office of the Assistant Secretary of Defense (Command, Control, Communications, and Intelligence), supports the DMS SPWG in developing DMS security policy and security architecture.

Position Paper, Gerald Short (TRW)

1. INTRODUCTION

The purpose of this paper is to document my views regarding the requirements for a security label for sensitive unclassified information. These requirements address operational, community, size, and mathematical concerns. The remainder of this paper covers the following topics: General Requirements, Detailed Label Composition,and Label Issues.

2. GENERAL REQUIREMENTS

Before evaluating the adequacy of the proposed standard, the general requirement sources must be identified to verify and validate the adequacy of the proposed standard.

a. First, the security label must be of sufficient size to provide the currently existing classification markings for security objects.

b. The security label has two purposes: to mark security objects and to be used in access control decisions. This latter use implies that there must be a minimum security value and a maximal security value. The minimum is dominated (in the mathematical sense) by all other values; while, the maximal value dominates all other values.

c. The label composition should provide for the following classes:

- o Identifies the specific hierarchical classifications.
- o Identifies a set of nonhierarchical entities - compartments in the classified world.
- o Accomodates future growth.

d. The label should be compatible with other existing standards that are evolving from various working groups.

3. Detailed Label Composition

This section provides detailed description of each of the above areas.

a. Classification field. The classification field must accommodate US concerns, NATO concerns, and privacy concerns. The hierarchical ordering suggested for message handling system under the Inter-Service Agency Automated Message Exchange was as follows:

TOP SECRET or COSMIC TOP SECRET
SECRET or NATO SECRET
CONFIDENTIAL or NATO CONFIDENTIAL
NATO RESTRICTED
UNCLASSIFIED CLEAR
UNCLASSIFIED EFTO

c. Additional security parameters. Additional security parameters are comprise any of the following:

o Caveats: The known phrases that comprise caveats are:

- DIRDIS used in conjunction with LOU or classified
- EXDIS means "Executive Distribution"
- EXCLUSIVE FOR or EXCLUSIVE-FOR will be followed by a proper name or title of the person receiving the message
- EYES ONLY is preceded by a combination of 2 letter foreign nations which can receive the message
- FOUO means "For Official Use Only"
- INSPECDIS means "INSPECTOR DISTRIBUTION"
- LIMDIS means "Limited Distribution"
- LOU means Limited Official Use"
- NOCONTRACT means "No Contractor"
- NODIS means "No Distribution"
- NOFORN means "No Foreign"
- ORCON means "Originator Control"
- PROPIN means "Proprietary Information"
- RELEASABLE TO ...
- RESDAT means "Restricted Data"
- RESTRICTED DATA
- SPECDIS used in conjunction with LOU or classified between director's office and addressees.
- STRADIS mean "State Distribution" only

4. LABEL ISSUES

There are at least the following consideration in the specification of a security label; hence any security label standard should be evaluated against these issues:

o Are there enough fields to handle all the levels, compartments, caveats, etc.

o To reduce the possibility of making a very large label to encompass all possibilities, should the label be a combination of bits and the characters.

o The rule should be explicitly stated for each component of the label. For example, the classification level will use integer arithmetic; while the non-hierarchical components will use boolean arithmetic. Further, there are at least two sets of rules for combining the non-hierarchical parts - Inclusive and Exclusive; that is compartments are ORed together and RELEASABLE are ANDed together.

o There are mathematical consideration so that "DOMINATES" function can apply and a lattice can be specified. A lattice requires that there be a maximal element and a minimal element. To aid this there should be an element known as "LATTICE HIGH" and another known as "LATTICE LOW". There should also be introduced the term as Site Accredited High (SAH) to replace the "System High" confusion with operating mode. SAH mean the highest classification and all the compartments that the site has been accredited for.

o Finally, there are some very practical consideration. There may be occasions when some data is received that should not be handled by the site. The term "NOSEND" or "ADMIN CONTROLLED" should be defined to information that cannot be sent out of the system or information that is controlled by the Security Administrator only"

"Communication Security Position Paper", Beverly Stein-Verbit (Department of Commerce Patent and Trademark Office)

Department of Commerce
Patent and Trademark Office (PTO)

Communication Security Position Paper
For the Workshop on Security Labels for Open Systems

Background - The Patent and Trademark Office processes and disseminates sensitive-unclassified information. The PTO uses communication facilities on a local, national, and global level.

PTO Host System Security - PTO data and word processing assets designated as sensitive-unclassified are subject to the Computer Security Act, OMB Circulars, DoC directives and guidance, and Federal government regulations. Within this context, sensitive-unclassified PTO systems are not required to establish and maintain NSA defined TCB B-1 label standards.

FIPS PUB. 146 - If GOSIP adopts TCB B-1 security labels, commercial product developers will strive to meet standards and guidance as published in the FIPS PUB 146. Organizations such as the PTO will necessarily assume the cost for products embedded with an inappropriate security measure.

PTO Position - Organizations with unclassified systems have a vested interest in the publication of GOSIP standards and guidelines. The PTO is encouraged by the invitational workshop as an opportunity to contribute to that effort.

"NIST Invitational Workshop on Security Labels", Leonard Tabacchi (Defense Communications Agency)

D C A INTEROFFICE MEMORANDUM

Date: 30-Apr-1990 06:34pm DST
From: Leonard Tabacchi
TABACCHIL
Dept: DCSO/DDEP/B615
Tel No: 703-285-5445

TO: NAZARIO@ECF.NCSL.NIST.GOV (REMOTE)

CC: Anthony Montemarano (MONTEMARANOA)
CC: James E. Mc Gowan (MCGOWANJ)

Subject: NIST Invitational Workshop on Security Labels for OSI

Someone from the DDN Project Office should be invited to attend this workshop on 30-31 May 90. The DDN PMO initiated the definition of the recently published Internet Protocol Security Option (IPSO). The DDN PM also administrates the assignment of Security Codes, Protection Authorities and compartments for use within DOD. We have a strong interest in having the IPSO adopted as the GOSIP standard for security labeling. DoD defined the IPSO to support the BLACKER communications security device which requires all data packets to be labeled. BLACKER currently supports users of the MILSTD protocols (TCP/IP) and is being modified to support users of GOSIP protocols and commercial X.25 connection oriented protocols. DoD has a strong interest in preserving the investment made in implementing the IPSO in the BLACKER system and in host computers required to interface to DDN through a BLACKER. Adoption of the IPSO labeling standard in GOSIP would also facilitate interoperability between OSI and MILSTD host computers duruing the transition to GOSIP. Use of the IPSO standard is not limited to DOD. IPSO is being used by the Department of Energy to label its data and will be used to label classified data transiting FTS2000.

Thanks, Len Tabacchi
Technical Manager, Defense Data Network

"An Approach for Labeling in Open, Heterogeneous, Distributed Systems", N. Vasudevan (IBM)

Position Paper for NIST Labeling Workshop

An Approach for Labeling in Open, Heterogeneous, Distributed Systems

N. Vasudevan

Secure Systems
Strategy & Planning
IBM Corporation

uunet!pyrdc!ibmsid!vasu
(301) 240-3716

An Approach For Labeling In Open, Heterogeneous, Distributed Systems

- *What Exists Today*

  - Commercial OS Products featuring Multi–Level Security exist

  - Labels for Objects related to Trusted Applications (Database, Windowing) can be designed based on the above

  - Limited Labeling "Standards" for Specific networks have been defined
    - IP Security Option for TCP/IP Networks

- *What We Need Next*

  - Labeling Standards for achieving **Inter–Operable Label Semantics** in Open, Heterogeneous Networks

An Approach For Labeling In Open, Heterogeneous, Distributed Systems

- *Framework for Discussion*
  - Labeling in End Systems
  - Labeling in Proprietary Networks
  - Labeling in Open, Heterogeneous Networks

Labeling in End Systems

- *Purpose of Labels in End Systems*

  - **Classification of Objects** based on the Sensitivity of Information

  - Assigning a **Range of Security Levels to Users** based on their Organizational Role

  - Statement of **MAC Policy** for the System

  - Enforcement of the **Reference Monitor** Concept

- *System Requirements Related to Labels*

  - Labels on **Objects** are **Trusted** (part of TCB)

  - Label **Functions** are **Trusted**

  - Subject Authorization related to Labels : **Authorization Functions and Data Base** are **Trusted**

Labeling in End Systems

- *Application of the Above*
  - Trusted OS
    - OS Subjects & Objects
  - Trusted Database
    - Database Subjects & Objects
  - Trusted Window System
    - Window System Subjects & Objects

Labeling in Proprietary Networks

- *Network Subjects & Objects Belong to a Single Security Domain*

  - NTCB consists of TCB Components which have **Mutual Trust**

  - **Label Semantics and Syntax is uniform** across the Network of TCBs (local "standard" for inter–operable labels)

- *System Requirements related to Labels based on the above*

  - No Requirement for Network to provide Authentication and Access Control

  - **Integrity of Labels** across the network is required

Labels in Open, Heterogeneous Networks

- *Properties*
  - Multiple Security Domains
    - No Mutual Trust
    - Diverse Label Semantics, Policies, etc.

- *Basic Issues*

  - **Need** for Information Exchange across Heterogeneous Security Domains

  - **Type** of Information to be exchanged and their **Sensitivities** (need for labeled information)

  - Can we list a **generic set** of Security Domains and their need for exchange of labeled information across domains?

    - Government
    - DoD
    - Commercial (proprietary, non–proprietary, ...)
    - Public Domain (legal, ...)

Labels in Open, Heterogeneous Networks

- *System Requirements*
  - Unified Label Semantics
    - Labeling Standards covering Clearances, Categories, etc.
  - Trusted Authorization of Subjects
    - via Trusted Inter–Domain Authorization Servers

Labels in Open, Heterogeneous Networks

- *Labeling Standards must allow for*
  - Future **Extensibility**
  - **Registration** of Label Standard for each Security Domain
    - Clearance & Categories
    - Syntax & Semantics

- *Label Management*
  - **Trusted Inter–Domain Services**
  - Performance Overhead
  - Transparency to the User

Labels in Open, Heterogeneous Networks

- *Basic Security Services needed*

  - **Authentication and Authorization Services** for Inter–Domain Access

  - **Access Control Service** by Inter–Domain Gateways (for controlled flow of Information based on Labels)

  - **Integrity Service** for Labels provided by Network Protocols

Labels in Open, Heterogeneous Networks

- *Extended Security Services needed*
  - **Confidentiality Service** for Labels
  - **Non–Repudiation Services**
    - Source
    - Destination
- *Approach for providing the above services:*
  - Use **Quality of Service** selection

"Position Paper Information Labels for NIST Security Labels for Open Systems Workshop", John P.L. Woodward (MITRE Corporation)

6

Position Paper on Information Labels for NIST Security Labels for Open Systems Workshop

John P. L. Woodward
The MITRE Corporation
5/15/90

Information labels were developed under the Defense Intelligence Agency/MITRE Compartmented Mode Workstation (CMW) project in an attempt to meet defense intelligence community data labeling requirements as well as requirements for controlling access to data based on classification and intelligence compartments. Although information labels were developed with the needs of the intelligence community in mind, they have utility outside intelligence and may even have an impact on the virus problem.

This paper describes the need for information labels and their advantages when used in conjunction with the more conventional sensitivity labels called for by the National Computer Security Center's *Trusted Computer System Evaluation Criteria* (DoD 5200.28-STD).

In trusted/secure computer systems that meet the B1 or higher requirements of DoD 5200.28-STD, sensitivity labels (SLs) are associated with subjects and objects for the purpose of implementing the system's mandatory access control (MAC) policy. MAC controls the access of subjects to objects based on the subject's classification and categories[1] and on the object's classification and categories. In studying how sensitivity labels could be used to satisfy the intelligence community needs for labeling human-readable data (e.g., printed by or otherwise exported from intelligence systems), four sensitivity label shortcomings were identified.

In reviewing these shortcomings, it is helpful to keep in mind that national policy controlling access to classified information seeks to maintain a balance between protecting unauthorized access to classified information and classifying information higher than it needs to be. As the following analysis will indicate, sensitivity labels can err too much on the side of controlling access at the (sometimes) expense of overclassification. Overclassification is a concern of the intelligence community, whose data--in the final analysis--is useful only to the extent that it can be released outside the intelligence community.

[1] Categories are the DoD 5200.28-STD term that is equivalent to the intelligence community's term compartments.

SENSITIVITY LABEL SHORTCOMINGS

Inability To Support Markings

First, sensitivity labels do not conveniently support *markings*. Markings are non-MAC-related information required to be associated with certain intelligence data products by national policy. Real-world examples of markings include NOFORN (not releasable to foreign nationals), NOCONTRACT (not releasable to contractors), PROPIN (proprietary information involved), ORCON (originator controlled), REL COUNTRY (releasable to COUNTRY), as well as many intelligence codewords. Although the National Computer Security Center's "Yellow Book" (*Guidance for Applying the Department of Defense Trusted Computer System Evaluation Criteria in Specific Environments*, CSC-STD-003-85) explicitly states that categories include some of the above examples, implementing these markings as categories in a sensitivity label has some drawbacks.

Although some markings imply access restriction, those that do all allow the originator of the data to grant exceptions to the restriction, e.g., to give NOCONTRACT data to selected contractors. Implementing these markings as categories in a sensitivity label does not allow for exceptions. Furthermore, some markings, including codewords, have *absolutely nothing* to do with access control, and are therefore inappropriately and inconveniently treated as categories. The discussion that follows further explores the difficulties associated with trying to force fit markings into sensitivity labels.

Force Fitting Access-Related Markings into Sensitivity Labels

It is possible to try to force fit access-related markings into sensitivity labels by treating them as categories. Consider the following example using the marking NOFORN. To treat NOFORN as a category:

1. Create a category named NOFORN

2. Give NOFORN data this category

3. Give U.S. citizens this category

4. Don't give non-U.S. citizens the category

There are two major ramifications of this treatment of NOFORN. First, it does not allow for exceptions. Even if the originator of some NOFORN data wants to grant access to a particular foreigner, the MAC enforced on the categories prevents it. It would be wrong to remove the NOFORN category from the data, thereby expanding its access to all foreigners. It would be similarly wrong to give the selected foreigner the NOFORN category, thereby granting him access to all NOFORN data.

The second ramification is that users must now change sensitivity labels more often to include or exclude the NOFORN marking on their products. If users find this inconvenient, they might always operate at the NOFORN level, thereby overclassifying information.

There is an alternative way of treating NOFORN with categories to allow for exceptions.

1. Create a category for NOFORN, and for each unique exception (e.g., NFxy for exception access to object x by foreigner y, etc.)

2. Give no-exception NOFORN data the NOFORN category

3. Give NOFORN data with exceptions the appropriate NFxy categories

4. Give all U. S. citizens the NOFORN and all NFxy categories (so they can read and write all types of NOFORN data)

5. Give foreigners the appropriate NFxy categories, for all values of x and y in the system

There are five major ramifications of this treatment of NOFORN. First, dynamic or frequent manual creation of categories is required. Second, many categories are required, possibly exhausting the available number. Third, the administration of such a system could get extremely cumbersome, with the security administrator having to keep track of all exception categories. Fourth, without extremely sophisticated software for creating human-readable sensitivity labels, sensitivity labels will look extremely non-standard and confusing to users. Fifth, the concern for users having to change their sensitivity labels often remains from the previous example, but is worse because of the larger number of label values.

The markings ORCON, NOCONTRACT, and PROPIN would be treated similarly to NOFORN, with the same classes of solutions and ramifications. A slightly different set of problems occur with release markings. Release markings, unlike true categories, expand access to data with which they are associated. Associating a true category with data restricts access to the data. Therefore, release markings cannot be treated directly as categories. They can be dealt with using categories in some inverse manner.

Consider the example depicted in the following figure, where there are three countries with users on a system. Categories could be assigned as suggested in the following tables, where NRCn means "not releasable to country n". Each object would be marked with categories indicating what countries to which the data is *not* releasable. Thus, generally releasable data has no categories. NOFORN data is marked as not releasable to all countries. Data releasable only to country 1 is marked as not releasable to the other two countries, etc. U.S. users are given all NRC categories so they can access all data. Citizens of a country are given the categories associated with all other countries. Even though this scheme may sound counter-intuitive, it works out if you study all of the dominance relationships. However, this scheme is extremely complicated, suffers from the same ramifications as the above schemes, and still does not allow for the exceptions inherent in the definition of the REL COUNTRY marking.

Type Of Data	Categories Assigned		
General Release	NONE		
NOFORN	NRC1	NRC2	NRC3
REL C1		NRC2	NRC3
REL C2	NRC1		NRC3
REL (C1, C2)			NRC3

Nationality Of User	Categories Assigned		
U.S.	NRC1	NRC2	NRC3
C1		NRC2	NRC3
C2	NRC1		NRC3
C3	NRC1	NRC2	

Force Fitting Non-Access-Related Markings into Sensitivity Labels

The final type of markings are those that are not access related at all, such as some intelligence codewords. Such markings must accurately be associated with certain types of information, but have absolutely nothing to do with access control. Treating these codewords as categories again forces users to change their sensitivity labels frequently to properly mark data, or to operate with all possible categories and therefore over-mark data.

Because of these difficulties in using categories for markings, we concluded that some other mechanism to handle markings was needed.

Forcing Predetermination of Output Product Classification

The second shortcoming is that sensitivity labels force predetermination of the classification of intelligence analysis products. Because sensitivity labels cannot dynamically change, users must choose a particular sensitivity label with which to work. If there are many categories being processed (which will be likely with categories used for markings, but true anyway for realistic intelligence analysis) and if users are producing analysis products that are combinations of potentially many different input products with different sensitivity labels (as is the case for many intelligence analysts), then users must guess in advance the proper sensitivity label for each output product. The problem is that if they guess too low they can't read all of the data they need to complete their analysis, forcing frequent upgrading of the product, and possibly logging in at successively higher levels. If they guess to high, the output product is overclassified and mismarked, and must be downgraded to be operationally useful. In this latter case, the practical problem that arises is how to accurately determine the proper sensitivity label if many items were combined.

Potential for Overclassification

The third, but related, shortcoming is that sensitivity labels can force the overclassification of information, again reducing its operational utility. For example, as illustrated below, if a SECRET subject makes a copy of a CONFIDENTIAL object, the copy of the object must be SECRET to avoid undesirable covert channels (see DoD 5200.28-STD for more information on covert channels).

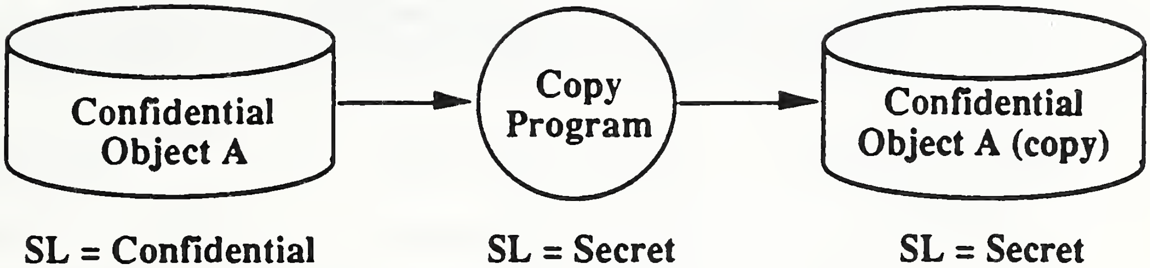

This problem is not strictly limited to programs that copy data, but rather to any application that reads some data objects and creates new ones.

Networking Systems With Difference Evaluation Classes (Degrees of Trust)

The fourth sensitivity label shortcoming is that, in a networking environment, it is unclear how a system should treat sensitivity labels less trustworthy than its own labels. For example, as depicted below, if a B1 system processing CONFIDENTIAL and SECRET sends data marked CONFIDENTIAL to an A1 system, can the A1 system believe the data is CONFIDENTIAL, or must it treat it as SECRET?

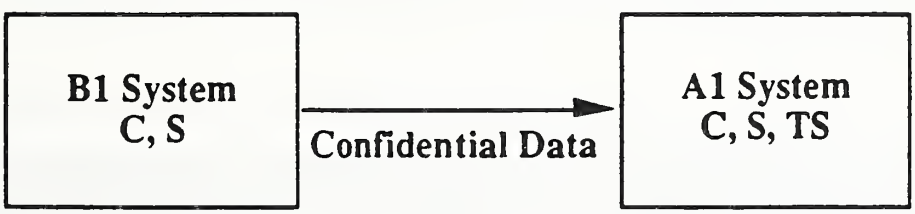

To retain its A1 trustworthiness, the A1 system must treat the data as SECRET, therefore possibly overclassifying the data and throwing away the CONFIDENTIAL label.

INFORMATION LABEL BACKGROUND

Information labels are one of the requirements defined in *Security Requirements for System High and Compartmented Mode Workstations*, DIA Document DDS-2600-5502-87. Information labels are a second label to be associated by a trusted computing base (TCB) with subjects and objects in computer systems, designed to work in conjunction with sensitivity labels. Whereas sensitivity labels are associated with subjects and objects themselves, information labels should be thought of as being associated with the *data* in subjects and objects—a key difference. Information labels contain a classification, categories (called compartments in DIA literature), and markings. The classification and categories components are analogous semantically to the classification and categories in sensitivity labels, but they are not used for access control. Markings are directly analogous to the markings described above.

Sensitivity labels *control* the flow of data, whereas information labels *track* the flow of data as it flows through the system. As shown in the example below, if data with an information label of SECRET NOFORN is combined with data with an information label of TOP SECRET, the resulting data's information label is automatically set to TOP SECRET NOFORN by the TCB.

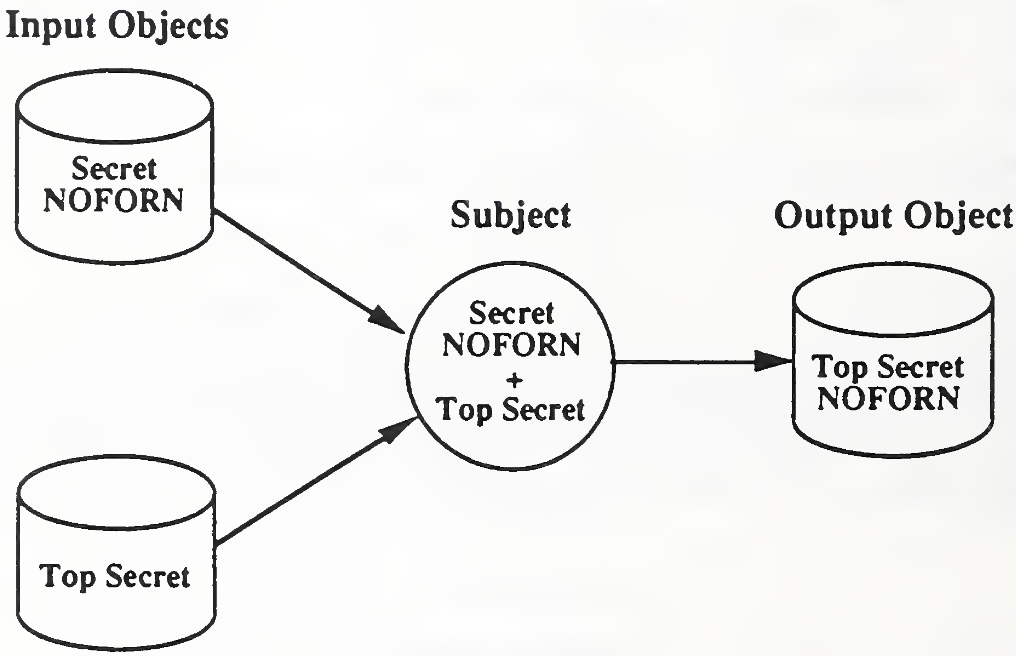

Information labels are intended to be used in conjunction with sensitivity labels in a trusted system, with the restriction that the classification and categories in the sensitivity label must dominate the classification and categories in the information label associated with the same object or subject. All subjects and objects should have both information and sensitivity labels. When an object is read by a subject (assuming that the MAC policy associated with the sensitivity labels allows the read), the information label of the object read is combined by the system with the information label of the subject, and the result becomes the new information label of the subject. Similarly when a subject writes an object (assuming that the MAC policy associated with the sensitivity labels allows the write), the information label of the subject is combined by the system with the information label of the object, and the result becomes the new information label of the object. The following table contrasts information and sensitivity labels.

	INFORMATION LABEL	SENSITIVITY LABEL
REPRESENTS:	Data in subjects and objects	Subjects and objects themselves
INITIALIZATION UPON CREATE:	UNCLASSIFIED (because there is no data present)	Inherited from creator (to prevent covert channel)
UPON CLEARING:	Reset to UNCLASSIFIED	Does not change
CHANGES:	Automatically on reads and writes	Never (except through extraordinary privilege)

HOW INFORMATION LABELS MITIGATE THE SHORTCOMINGS

Inability to Support Markings

Information labels directly solve the inability of sensitivity labels to support markings in that they explicitly contain markings and do not force their use for MAC. Furthermore, they facilitate the proper marking of data by allowing data to be properly marked when it is entered into the system, but then automatically marks output products based on the combination of the information labels of the objects that went into the output product.

Forcing Predetermination of Output Product Classification

Information labels mitigate the problem relating to predetermining the output product sensitivity label, because they automatically compute the proper information label of the output product, allowing the analyst to perform the analysis at his/her maximum sensitivity label, thereby allowing access to all needed data. The figure below shows a typical example of such an analysis. By operating at his/her clearance (Top Secret with compartments A, B, and C), the user can access all data potentially available to perform the analysis. Each potential type of data is shown with a proper information label. The information label includes a codeword that identifies the compartment in which the data is protected (CodewordA for compartment A, etc.). The output product constructed by the user is automatically protected with a sensitivity label of TS A B C, because the user could put data up to that level in the output object. However, when the analysis is complete, the user finds that data from all compartments was not needed in performing the analysis (no data from compartment C was needed). The user can tell this from the information label automatically computed for the output product (because CodewordC does not appear in the information label). The user can then downgrade the sensitivity label as necessary using the information label as a guide (i.e., the information label is used as a downgrading hypothesis). Note that the user did not have to guess in advance the proper sensitivity label at which to work to complete the analysis.

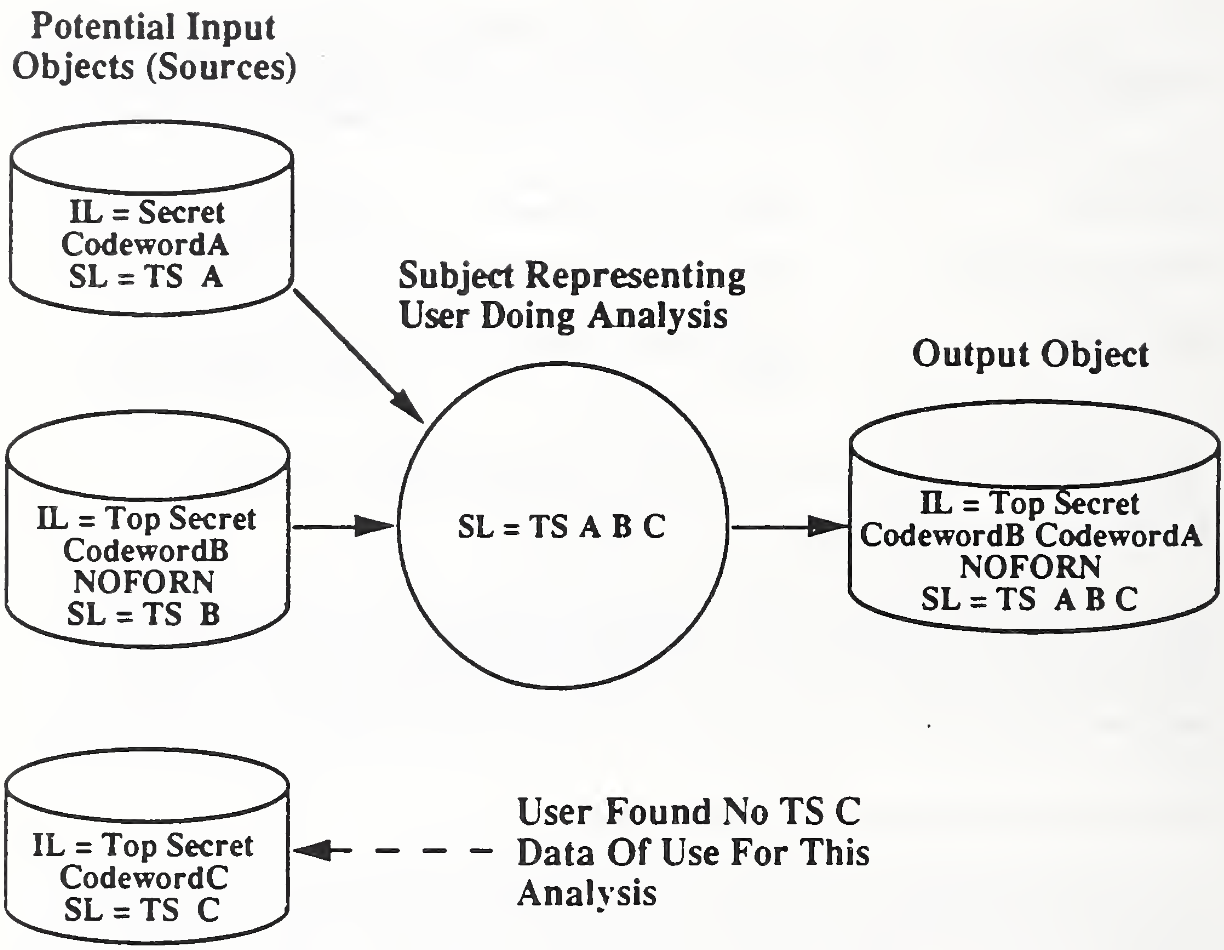

Potential for Overclassification

As shown below, information labels can mitigate the overclassification problem mentioned above, in that although the copy of the CONFIDENTIAL object by the SECRET process does indeed cause the new object's sensitivity label to be SECRET, the object's information label would be CONFIDENTIAL, and would therefore more accurately represent the classification of the data in the object. If operationally needed, the sensitivity label of the copied object could be downgraded, again using the information label as the downgrading hypothesis.

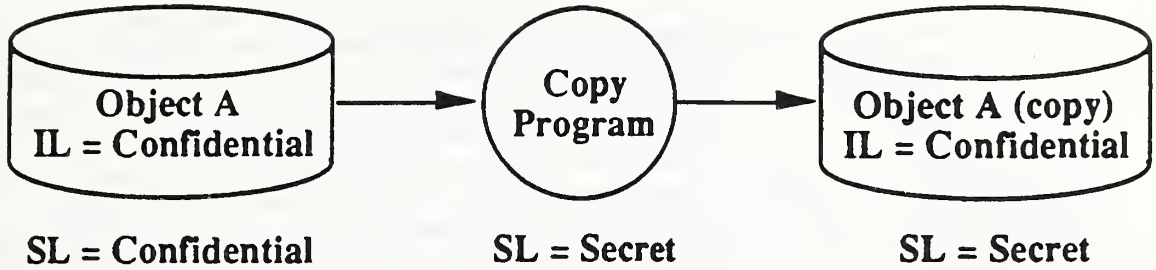

Networking Systems With Difference Evaluation Classes (Degrees of Trust)

Information labels address the problem that arises with networking systems with different degrees of trust because they can be used to store less trustworthy sensitivity labels. Therefore, as shown below in a modification of the original example, the A1 system would, upon receiving data with a CONFIDENTIAL sensitivity label from the B1 system, set the sensitivity label of the received data to SECRET to retain its A1 trustworthiness, yet retain the original information label of the received data as CONFIDENTIAL, such that CONFIDENTIAL could later be used as a downgrading hypothesis.

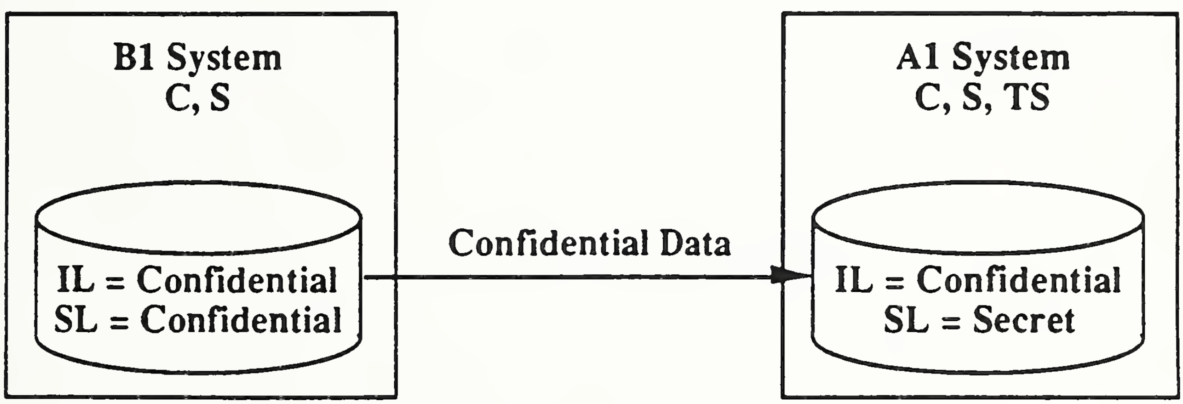

HOW INFORMATION LABELS CAN MITIGATE THE VIRUS PROBLEM

Finally, information labels can be used to mitigate the virus problem, not by preventing viruses, but by detecting them. For example, if suspect programs (e.g., those pulled off bulletin boards) are given information labels with unique markings, then any files they surreptitiously modify will automatically inherit the unique marking. The system can then be regularly scanned for the occurrence of files with these markings.

CONCLUSION

In conclusion, this position paper, because of its brevity, barely scratches the surface of information labels. They are being implemented in commercial workstations under the CMW program by DEC, Harris, IBM, SUN Microsystems, and SecureWare. DIA standards for translating information labels and sensitivity labels between human-readable and internal encoded formats, as well as for specifying which classifications, compartments, and markings are being processed on each system, are documented in *Compartmented Mode Workstation Labeling: Source Code and User Interface Guidelines*, DIA Document DDS-2600-6215-89. Information labels are also present in the workstation being fielded by Honeywell for the World-Wide Military Command and Control Systems (WWMCCS). Information labels, though designed for intelligence community needs, have additional utility outside the intelligence community—possibly even in the commercial sector. Because they track the flow of data rather than prevent the flow, they can be useful in reconstructing some system activity after abnormal behavior is discovered, or they can be used to detect viruses.

"Background on Extended Security Options and Compartments in IP/CLNP Labels" (Presentation Slides), John P.L. Woodward (MITRE Corporation)

BACKGROUND

- o Why are there ESO's (133's)?
  - Because across-the-board agreement on format of non-classification could not be reached
  - Break the problem down into smaller pieces (1 per ESO)
- o The Problem with ESO's?
  - Allows for many encodings of the <u>same</u> semantical information

INTELLIGENCE COMMUNITY STANDARDS

- CSESO - Common SCI ESO
  - Compartments commonly exchanged among agencies
  - "What we could agree on"
- RMESO - Release Markings ESO
  - For countries potentially internetworked
- SDESO - Supplemental Data ESO
  - "Other Stuff"
    - Subcompartments
    - Handling Restrictions
- SIESO - DNSIX (DoDIIS [DOD Intelligence Information Systems] Security for Information Exchange) Session ID

HOW ESO's ARE PROCESSED

- o CSESO & RMESO - used for MAC, bit encoded
  - Why two ESO's?
    - o Default for CSESO if missing: all 0's
    - o Default for RMESO if missing: all 1's
  - Both what DNSIX calls Network Level ESO's (NLESO)
- o SDESO - integer encoded, processing potentially unique per value
- o SIESO - integer encoded, all datagrams without known SIESO's rejected

DNSIX NETWORK LEVEL MODULE

- o Checks datagrams going in or out of hosts or IP Gateway Interfaces
- o Checks:
  - That there is one BSO
  - BSO classification valid (one of the 8)
  - PAF value found in PAF table (exact or dominates match depending on table entry)
  - There is at most one ESO per type (Source; prot. auth)
  - Each ESO is in NLESO table, or Auxillary ESO table
  - The network level (BSO classification and category bits from all NLESO's) are found in accreditation range table (exact or dominates match, depending on table entry) (all dominates for incoming datagrams

NLESO TABLE

- o Type (Source, Protection Authority)
- o Max size (# Categories)
- o Default 1 or 0
  - CSESO 0
  - RMESO 1

Attendees List - NIST Invitational Workshop

ATTENDEES LIST - NIST Invitational Workshop

Branstad Dennis
NIST

Brown C. Douglas
Sandia National Laboratories

Brozyna David
Department of Defense

Carnahan Lisa
NIST

Chambers Douglas
U.S. Patent & Trademark Office

Chizmadia David
National Computer Security Center

Eshgh Ali
SSDS, Inc.

Fadden Fran
Digital Equipment Corporation

Gitlin Julian
UCEC

Housley Russell
Xerox Special Information Systems

Hovey Richard
Digital Equipment Corporation

Katzke Stuart
NIST

Kutz Bill
National Computer Security Center

Linn John
Digital Equipment Corporation

Maimone Bill
Oracle Corporation

Mellinger Philip
MITRE Corporation

Migues Sammy
Xerox

Nazario Noel
NIST

Pope Nick
Logica (U.K.)

Rogers George
Intelligence Community Staff

Rogers Kristina
MITRE Corporation

Rosenthal Robert
NIST

Schmitt Warren
Sears Technology Services

Shirey Robert W.
MITRE Corporation

Short Gerald
TRW

Stein-Verbit Beverly L.
U.S. Department of Commerce

Steinauer Dennis
NIST

Sudduth Larry
National Capitol Systems, Inc.

Tannert Robert A.
Department of Energy HDQTRS.

Tarquinio Sally
MITRE Corporation

Troy Gene
NIST

Vasudevan N.
IBM

Walters Dale
NIST

Wood Lauren
Allied Signal Aerospace Co.

Woodward John
MITRE Corporation

NIST-114A (REV. 3-89)	U.S. DEPARTMENT OF COMMERCE NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY	1. PUBLICATION OR REPORT NUMBER NISTIR 4362
	BIBLIOGRAPHIC DATA SHEET	2. PERFORMING ORGANIZATION REPORT NUMBER
		3. PUBLICATION DATE JUNE 1990

4. TITLE AND SUBTITLE

Security Labels for Open Systems - An Invitational Workshop

5. AUTHOR(S)

Noel A. Nazario, Chairman

6. PERFORMING ORGANIZATION (IF JOINT OR OTHER THAN NIST, SEE INSTRUCTIONS)	7. CONTRACT/GRANT NUMBER
U.S. DEPARTMENT OF COMMERCE NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY GAITHERSBURG, MD 20899	8. TYPE OF REPORT AND PERIOD COVERED

9. SPONSORING ORGANIZATION NAME AND COMPLETE ADDRESS (STREET, CITY, STATE, ZIP)

National Institute of Standards and Technology
Bldg. 225/Rm. A216
Gaithersburg, MD 20899

10. SUPPLEMENTARY NOTES

[ ] DOCUMENT DESCRIBES A COMPUTER PROGRAM; SF-185, FIPS SOFTWARE SUMMARY, IS ATTACHED.

11. ABSTRACT (A 200-WORD OR LESS FACTUAL SUMMARY OF MOST SIGNIFICANT INFORMATION. IF DOCUMENT INCLUDES A SIGNIFICANT BIBLIOGRAPHY OR LITERATURE SURVEY, MENTION IT HERE.)

On May 30 and 31, 1990 the Protocol Security Group at NIST held a Workshop on Security Labels. Thirty-Five representatives from the U.S. Government, industry, and the United Kingdom gathered for two days to discuss Security Labels for open systems. The discussion went from the generalities of labels in "end systems" to the more specific issue of labels in secure Open Systems Interconnection (OSI). The information shared during the two days discussion is documented in these proceedings.

12. KEY WORDS (6 TO 12 ENTRIES; ALPHABETICAL ORDER; CAPITALIZE ONLY PROPER NAMES; AND SEPARATE KEY WORDS BY SEMICOLONS)

Government open systems interconnection profiles; Integrity; Network; Open systems interconnection; security labels; trust

13. AVAILABILITY	14. NUMBER OF PRINTED PAGES
[X] UNLIMITED	196
[ ] FOR OFFICIAL DISTRIBUTION. DO NOT RELEASE TO NATIONAL TECHNICAL INFORMATION SERVICE (NTIS).	
[ ] ORDER FROM SUPERINTENDENT OF DOCUMENTS, U.S. GOVERNMENT PRINTING OFFICE, WASHINGTON, DC 20402.	15. PRICE A09
[X] ORDER FROM NATIONAL TECHNICAL INFORMATION SERVICE (NTIS), SPRINGFIELD, VA 22161.	

ELECTRONIC FORM