Network Working Group
INTERNET-DRAFT

Thomas C. Bartee (IDA)
Nelson W. Alvarez (DISA)
C. Dale. Nunley (DoD)
June 1997

Internet Security Label (ISL)
<draft-tbnadn-sec-label-00.txt>

Status of this Memo

Abstract

This document describes the Internet Security Label (ISL). ISL provides a mechanism for encoding security (sensitivity) parameters. The ISL is intended to be a layer-independent security mechanism. It can be used with all current versions of the Internet Protocol (IP), including IPv4 and IPv6 as well as the IP Security Protocols (IPSEC), the encapsulating Security Payload (ESP) and the Authentication Header (AH). Other protocols which use a security label can also use the ISL encoding standard including IPX.

Table of Contents

1. Introduction 3
1.1 Overview 5
1.2 Requirements Terminology 5
1.3 Technical Definitions 6

2. General Description 10
2.1 Basic 10
2.2 General Tag Format 11
2.3 Tags 12
3. Access and Routing Control 19
3.1 Clearance Considerations 19
3.2 Error Processing 20
3.3 Example Tag Procedures 21
4. Security Considerations 22
5. References 23

1. Introduction

A security label is an indication of the protection requirements for information (including programs and data)

imposed by a security policy. Although security labels have often been restrictively limited to "sensitivity" of information related to access control, other aspects of security policy (e.g., confidentiality requirements, integrity requirements, non-repudiation requirements) may also need to be represented. Familiar examples of labels indicating sensitivity are: the "Company Confidential" marking used by many corporations to indicate the material so marked is not to be released to other than corporation (company) employees (unless corporation management makes such a release.); "Corporate Financial" is often used to restrict release to only those working in the financial area; "Medical Records" is used to enforce privacy laws concerning employee (or patient) private medical information, etc. Agencies of the Government have similar markings which restrict data release to other than selected groups and military establishments generally have complex marking systems to control access to sensitive information (this includes NATO.) Inter-Government organizations such as IMF, World Bank, etc. also have marking systems.

There may be several markings associated with particular data and we find "AMEX Company Confidential", "Release to Film Development Only" on a single document or protocol entity. In this RFC, the collection of all the security markings associated with a protocol entity is called a security label. Internet RFCs also often use the term sensitivity label [2].

Security labels convey information used by protocol entities, operating systems, and applications to determine how to protect information in open systems. Information in a security label can be used to control access, specify protective measures, and determine additional handling restrictions required by a security policy.

The syntactic constructs for conveying security information in this standard are used along with the semantics provided by the security authority which establishes security policy for the information exchanged. It is anticipated that each security domain will have a registration authority whose responsibility will be to register security-related technical objects for that domain. These objects could include security policies, certificate policies, security labels, and others. The registration authority will ensure that the domain

number/object identifier is unique for each technical object. The security domain registration authority must ensure that registration applications contain the appropriate information such as name, address (physical and Email), telephone number, person to contact, releasability of the registration information, a proposed alpha-numeric name for the domain, domain description, and any other information required by the security domain registration authority. Dissemination of the registration information will be at the discretion of the domain security policy administrators, however it must be available to all authorized components that communicate within that domain.

There are two types of markings commonly used in security labels which are called "restrictive markings" and "release markings." There are also "hierarchical" or "level markings" which form a special subset of restrictive markings.

When access to information is to be limited to only those who qualify by being included in "every" marking in a section of a label the markings are "restrictive markings." An example is where a Corporation wants to release information to only those who work for the corporation, and are in an engineering development program, and who are in the financial department. The markings might then be "Genzyme"; "engineering"; "financial" where "Genzyme" is a restrictive marking, "engineering" is a restrictive marking and "financial" is a restrictive marking.

"Release markings" are also used to control information dispersal and these include such markings as "Release to Amgen Biotech", "Release to First Genome UK Division." etc. In order to qualify for receival of information when several of these markings are used in a label only one (or more) of the markings is required. For the immediately preceding two markings a person would qualify if he/she worked in Amgen Biotech "or" was with First Genome in the UK.

"Hierarchical Markings" are a subset of restrictive markings which are ordered in that information (or someone) in a high level category is always in all of the lower level categories. The most familiar instances are the "unclassified-confidential-secret-top secret" markings used by several Governments where information which is secret, for instance, is considered more sensitive than information which is confidential, or unclassified. Those who have been categorized by a Government as able to accept secret information would then also be

qualified to receive lower level information (generally with the qualification that they have some reason to know the information.)

In order to transmit the sensitivity markings in a communications system and to control access to the associated information an encoding mechanism can be used. The encoding mechanism is the subject of this document. This encoding standard permits separation of communities (domains) such as Corporations, Government Agencies, etc. and encoded marking interpretations are unique within communities (domains.) The encoded markings can be used for access control before transmission, on the network (in routers, for example), and at receivers or during processing. This document assumes some familiarity with the TCP/IP headers and with the Internet "IP Security Architecture" document [2] which provides background information.

1.1 Overview

The Internet Security Label provides the ability to control and communicate security information for data such as printed documents, display windows and database records. The ISL uses standard computer encoding formats. The ISL enables users to make efficient (compact) encodings. Processing is also fast in order to facilitate usage in routers, gateways, etc. Tables can be used to provide mappings between ISL encodings and operating systems and/or network specific encodings. As is the case for Tunnel-mode ESP and Transportation-mode ESP, the encoded markings can be placed in the encapsulating security payload or the IP datagram prior to the transport-layer protocol header (TCP, UDP, ICMP) respectively. It can be used in AH in an unencrypted IP packet, after the IP header and before the ESP header (for transport-mode ESP) or in the encrypted tunnel-mode ESP.

1.2 Requirements Terminology

In this document, the words that are used to define the significance of each particular requirement are usually capitalized. These words are:

- MUST
  This word or the adjective "REQUIRED" means that the item is an absolute requirement of the specification.

- SHOULD

This word or the adjective "RECOMMENDED" means that there might exist valid reasons in particular circumstances to ignore this item, but the full implications should be understood and the case carefully weighed before taking a different course.

- MAY
This word or the adjective "OPTIONAL" means that this item is truly optional. One vendor might choose to include the item because a particular marketplace requires it or because it enhances the product, for example; another vendor may omit the same item.

1.3 Technical Definitions

This section provides a few basic definitions that are applicable to this document. Other documents provide more definitions and background information. ([3] - [6].)

Authentication (of Data Origin):
A security property that ensures that the origin of received data is, in fact, the claimed sender. Usually bundled with integrity services, especially connectionless integrity.

Bit Order:
The ISL assumes a standard ordering of bits as they are transmitted over a network. Bits within bytes are transmitted from most significant bit (MSB) to least significant bit (LSB).

Confidentiality:
A security property that ensures that communicated data is disclosed to intended recipients, but is not disclosed to other, unauthorized parties. Traffic flow confidentiality extends this service to externally visible characteristics of communication, e.g., source and destination identifiers, message length, or frequency of communication. (See also traffic analysis, below.)

Destination system:
This is the information system that is identified by the

destination address in the protocol data unit header. This system will be the one to receive the protocol data unit and pass it to an upper layer protocol.

DOI:
A collection of systems that share a same set of security

policies and a common interpretation of security attributes form a Domain of Interpretation (DOI).

DOI authority:
The DOI authority is the organization that has obtained a DOI identifier. The authority is responsible for defining and requesting registration for and distributing the DOI mapping.

DOI Identifier:
The DOI Identifier is the number used in the ISL header to uniquely represent a DOI.

Encryption:
A mechanism commonly used to provide confidentiality.

End system:
This refers to either the source or destination system.

Intermediate System:
A system performing functions of the lower three layers of the OSI Reference Model, commonly thought of as routing data for end systems.

Integrity (Connectionless):
A security property that ensures data is transmitted from source to destination without undetected alteration. If the order of transmitted data also is ensured, the service is termed connection-oriented integrity. The term anti-replay refers to a form of connection-oriented integrity designed to detect and reject duplicated or very old data units.

Label range:
This is a pair of security labels which bound the security labels a subject may access. It is assumed that all objects whose labels fall within this range, inclusive, are accessible by the subject.

MLS:

Multi-Level Security (MLS) is the practice of giving end systems or network resources a sensitivity label and restricting access to those resources based on a users clearance label range.

Network byte order:
The ISL assumes the most significant byte/octet is transmitted first.

Non-repudiation:

A security property that ensures that a participant in a communication cannot later deny having participated in the communication. This property may apply to either the sender or the recipient of communicated data, or both.

Object:
An object is a network resource to which access by a subject must be controlled in accordance with the local security policy. Examples of objects for networks are: protocol data units, applications, configuration parameters, connections, and networks.

Protocol Data Unit:
A unit of data specified in a protocol and consisting of protocol information and, possibly, user data.

Release marking:
A release marking provides a list of authorized subjects who may access the associated object. The release marking is made up of release categories.

Release category:
The release category represents a subject or group of subjects that may access an object.

SPI:
Acronym for "Security Parameters Index." The combination of an SPI and a destination address uniquely identifies a simplex security association (SA, see below). The SPI is carried in IPsec protocols to select the set of parameters bound to an SA; thus an SPI is generally viewed as an opaque bit string. However, the creator of an SA may choose to interpret the bits in an SPI to facilitate local processing.

Security Association (SA):
A simplex (uni-directional) logical connection, created for security purposes. All traffic traversing an SA is subjected to the same security processing at the transmitter and receiver. In Ipv4 and Ipv6 security (IPsec), an SA is a network layer abstraction enforced through the use of AH or ESP.

Security attribute:
A security-related quality of an object.

Security Gateway:
A system that acts as the communication interface between

untrusted external, networks and internal (trusted) hosts and subnetworks. The internal subnets and hosts served by a security gateway are presumed to be trusted by virtue of sharing a common, local, security administration. (See "Trusted Subnetwork" below.) In the IPsec context, a security gateway is a point at which AH and/or ESP is implemented in order to serve a set of internal hosts, providing security services for these hosts when they communicate with external hosts also employing IPsec (either directly or via another security gateway).

Security domain:
A collection of entities to which applies a single security policy managed by a single authority.

Sensitivity category:
A sensitivity category is a security attribute that describes in absolute terms a protection requirement.

Security Policy Identifier:
An identifier that may be used to identify the security policy enforced.

Sensitivity level:
A security attribute that indicates a required level of confidentiality protection according to a predefined protection hierarchy. The level is hierarchical in ascending order, meaning that level N represents greater sensitivity than level N-1.

Source system:
This is the information system that originated the protocol data unit with the ISL label.

Subject:
The subject is an active entity that requests access to a particular object. Examples of subjects are hosts, network, computer processes, and users. A subject can also be an object.

Trusted Subnetwork:
A subnetwork containing hosts and routers that trust each other not to engage in active or passive attacks and trust the underlying communications channel (e.g., an Ethernet) is not being attacked.

Vector, binary valued:

An n-component binary-valued vector A = (a(0),a(1),...,a(n)) is an ordered list of n binary values a(i) = 0 or 1.

Vector sum:
The sum of two n-component binary-valued vectors A = (a(0),a(1),...,a(n)) and B = (b(0),b(1),...,b(n)) is A + B = (a(0)+b(0),a(1)+b(1),...,a(n)+b(n)) where 0 + 0 = 0 and 0 + 1 = 1 + 0 = 1 + 1 = 1.

Vector Product:
The product of two n-component binary-valued vectors A = (a(0),a(1),...,a(n)), and B = (b(0),b(1),...,b(n)) is A * B = (a(1) b(1),a(2) b(2),...,a(n) b(n)) where 0 * 0 = 0 * 1 = 1 * 0 = 0 and 1 * 1 = 1.

2. General Description

The ISL allows the attachment of specific security attributes to data in a protocol data unit. The security attributes can be used to perform security decisions. ISL can support a large set of security domains and policies with differing interpretations of security attributes. An extendible format allows for multiple sets of security attributes as well as the addition of new attribute types in the future. This document defines the basic formats and gives

example processing procedures..

2.1 Basic Format

The basic format of the ISL is shown below. A fixed format header is followed by a variable length tag section.

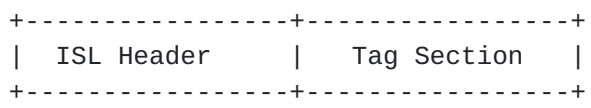

Figure 1: General ISL Format

A single ISL may include multiple tags.

2.1.1 Header

The ISL header identifies the ISL and contains the Domain of Interpretation (DOI) Identifier which is used to interpret the tag section.

2.1.2 Format

The format for the ISL header is shown in Fig. 2.

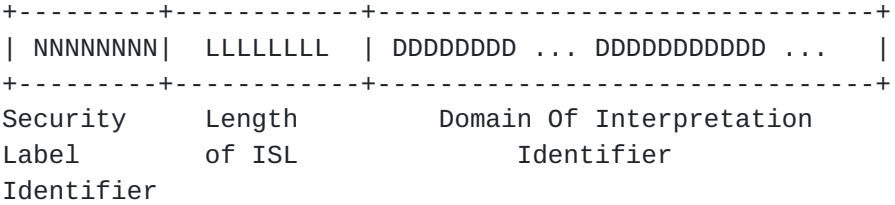

Figure 2: ISL Header Format

2.1.3 Security Label Identifier

The first field is a one octet security label identifier field.

2.1.4 Length of ISL

The one octet ISL length field represents the length in

octets of the entire ISL including all tags and the ISL security label identifier and length fields.

2.1.5 Domain of Interpretation (DOI) Identifier

The DOI Identifier is 4 octets in length and stored in network byte order. The security attributes contained in the tag section will have meaning to systems within the same security domain as specified by the DOI.

2.2 General Tag Format

Tags are independent data elements within an ISL that convey security attributes. An ISL will include 0 or more tags in any order. The purpose of tags is to provide an extensible method to pass security attributes using predefined formats and relating to a general security policy.

2.2.1 Format

The standard format for an ISL tag is shown below.

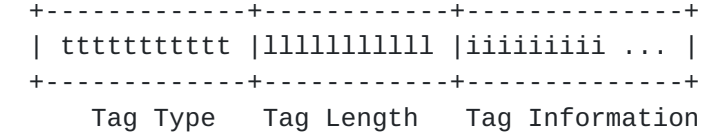

Figure 3: General Tag Format

The tag type is one octet in length and is used to identify the specific format and processing procedures associated with the tag information field. The section below provides more information on tag types. The tag length is one octet in length and gives the total octets in the tag including the tag type and length fields.

2.2.2 Tag Type

The tag type is a number between 0 and 255. Tag types 0 through 127 are used for standard tag definitions. For these standard tags the tag type number alone will identify the

format for the tag information field. The DOI then determines the semantics for a given tag. The ISL defines standard tag types 1, 2, 5, 6 and 7. Tag types 0, 3, and 4, and 8 through 127 are currently reserved for future use. Tag types 128 through 255 can be defined by the DOI authority. This makes it possible to use a unique tag specification for such domain specific things as human readable time stamps, policy identifiers and privacy marks. This provides "free form" tags which may not be registered, although registration with IANA is recommended.

2.3 Tags

The tags defined below represent three different ways to format a sensitivity label. Each of them store a sensitivity hierarchical level in a one octet field.

2.3.1 Tag Type 1

This is referred to as the "bit-mapped" tag type. The format of this tag type is as follows:

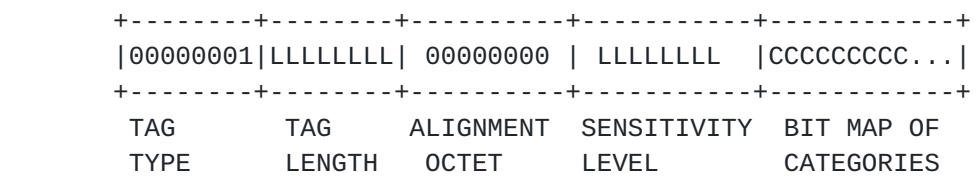

Figure 4. Tag Type 1 Format

2.3.1.1 Tag Type

This field is 1 octet in length and has a value of 1.

2.3.1.2 Tag Length

This field is 1 octet in length. It gives the total length of the tag in octets including the type and length

fields.

2.3.1.3 Alignment Octet

This field is 1 octet in length and always has the value of 0.

2.3.1.4 Sensitivity Level

This field is 1 octet in length. Its value is a binary number with value from 0 to 255 decimal. The values are ordered with 0 being the minimum value and 255 representing the maximum value. These values represent sensitivity levels.

2.3.1.5 Bit Map of Categories

The length of this field is variable and ranges from 0 to 30 octets. This provides representation of sensitivity categories 0 to 239. The ordering of the bits is left to right or MSB to LSB. For example category 0 is represented by the most significant bit of the first byte and category 15 is represented by the least significant bit of the second byte. Figure 5 graphically shows this ordering. Bit N is binary 1 if category N is part of the label for the protocol data unit, and bit N is binary 0 if category N is not part of the label. Minimal encoding should be used resulting in no trailing zero octets in the category bit map. That is, the final right octet in a bit map in a transmitted ISL should contain at least a single 1.

	octet 0	octet 1	octet 2	octet 3	octet 4
	XXXXXXXX	XXXXXXXX	XXXXXXXX	XXXXXXXX	XXXXXXXX
bit	01234567	89111111	11112222	22222233	33333333
number		012345	67890123	45678901	23456789

Figure 5. Ordering of Bits in Tag 1 Bit Map

A bit map is a binary-valued vector V = (v(0),v(1),...,v(8n-1)) where v(i) = 0 or 1 and where n is the number of octets in the vector. Each category to be represented in the vector for a specific DOI is assigned a position in the vector corresponding to an i in a specific v(i). If the value of v(i) is a 1 the category is in the ISL.

If the value of v(i) is 0 then the category is not in the ISL. For example, if v(2) is assigned the category "Kodak" and v(3) the category "Fuji" then if v(2)v(3) = 01 the ISL bit map indicates the marking "Fuji" (and not "Kodak.") The final rightmost octet in a transmitted ISL category bit map should contain at least one v(i) = 1.

2.3.2 Tag Type 2

This is referred to as the "enumerated" tag type. It can be used to describe large sets of sensitivity categories. The format of this tag type is as follows:

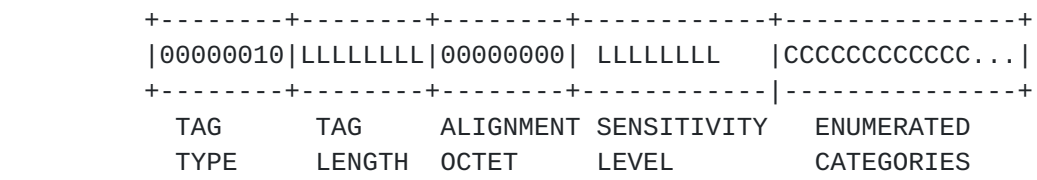

Figure 6. Tag Type 2 Format

2.3.2.1 Tag Type

This field is one octet in length and has a value of 2.

2.3.2.2 Tag Length

This field is 1 octet in length. It gives the total length of the tag type including the type and length fields.

2.3.2.3 Alignment Octet

This field is 1 octet in length and always has the value

of 0.

2.3.2.4 Sensitivity Level

This field is 1 octet in length. Its value is from 0 to 255. The values are ordered with 0 being the minimum value and 255 representing the maximum value.

2.3.2.5 Enumerated Categories

In this tag, a category is represented by a numerical value rather than by a position within a bit map. The length of the enumerated category field is 0 to 251 octets. The length of each category number is 2 octets. Valid values for categories are 0 to 65534 decimal. Category 65535 is not a valid category value.

Since each category to be represented by a specific DOI is assigned a 16 binary-bit value, a table can be made of the assigned values. For example, if the category "Sylvania" is assigned the value 0000010000010010 and the category "Samson" the value 0000010000010011 a section of this table would be

0000010000010010 = Sylvania
0000010000010011 = Samson

Note that alphanumeric codes can be used to assign values to the 2 octet ISL enumerated category numbers. For example, if, for a specific DOI, SV is used for Sylvania and SA for Samson, the ANSI-ASCII codes for A, S, V are 10000011, 10100111, and 10101101 (with odd parity) and so the assignment would be

1010011110101101 = SV = Sylvania
1010011110000011 = SA = Samson

Each 2-octet number is a 16 binary-bit vector V = (v(0),v(1),...,v(15)) where each v(i) = 0 or 1. Each ISL tag type 2 then contains a list of vectors each representing a category. Each category associated with an ISL is then represented by a vector in the ISL list.

2.3.3 Tag Type 5

This is referred to as the "range" tag type. It is used to represent labels where all categories in a range, or set of ranges, are included in the sensitivity label. The format of this tag type is as follows:

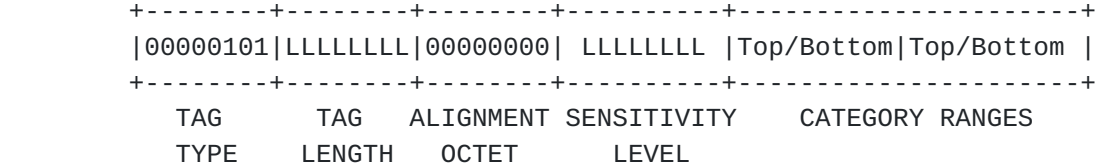

Figure 7. Tag Type 5 Format

2.3.3.1 Tag Type

This field is one octet in length and has a value of 5.

2.3.3.2 Tag Length

This field is one octet in length. It gives the total length of the tag type including the type and length fields.

2.3.3.3 Alignment Octet

This field is one octet in length and always has the value of 0.

2.3.3.4 Sensitivity Level

This field is one octet in length. Its value is from 0 to 255. The values are ordered with 0 being the minimum value and 255 representing the maximum value.

2.3.3.5 Category Ranges

A category range is a 4-octet field comprised of the 2-octet index of the highest-numbered category followed by the 2 octet index of the lowest-numbered category. These range endpoints are inclusive within the range of categories. All categories within a range are included in the sensitivity label. This tag may contain a maximum of 7 category pairs. Figure 7 shows two categories pairs. The ranges must be non-overlapping and be listed in descending order. Valid values for categories are 0 to 65534. Category 65535 is not a valid category value.

2.3.4 Tag Type 6

This is referred to as the "release markings" tag type. The format of this tag type is as follows:

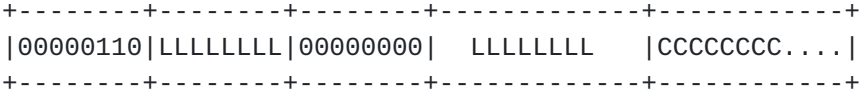

TAG TYPE	TAG LENGTH	ALIGNMENT OCTET	SENSITIVITY LEVEL	BIT MAP OF RELEASE CATEGORIES

Figure 8. Tag Type 6 Format

2.3.4.1 Tag Type

This field is one octet in length and has a value of 6.

2.3.4.2 Tag Length

This field is one octet in length. It gives the total length in octets of the tag type including the type and length fields.

2.3.4.3 Alignment Octet

This field is one octet in length and always has the value 0.

2.3.4.4 Sensitivity Level

This field is one octet in length. Its value is from 0 to 255. The values are ordered with 0 being the minimum value and 255 representing the maximum value.

2.3.4.5 Bit Map of Release Categories

The length of this field is from 0 to 251 octets. The bit map has one bit for each release category. All combinations are possible. The ordering of the bits is left-to-right or MSB to LSB. For example category 0 is represented by the most significant bit of the first byte and category 15 is represented by the least significant bit of the second byte. Figure 5 graphically shows this ordering. Minimal encoding should be used resulting in no trailing all zeros octets in the release category bit map.

The bit encoding used for release categories is the same as in the restrictive category encoding where a binary 1 means that the category is included in the label.

The release category bit map in an ISL is a vector V =

(v(0),v(1),...,v(8n-1)) where each v(i) is a 0 or 1 and where n is the number of octets in the vector. Each category is associated with a bit position in the vector. If for a particular DOI, AMGEN is assigned v(3) and BIOGEN v(4), then v(3) = 1, v(4) = 0 would mean the protocol data unit's contents are released to AMGEN and not to BIOGEN. The vector transmitted would then be 00010000 if the protocol data was to be released only to AMGEN. The final octet in a transmitted ISL release category bit map should not be all 0s.

2.3.5 Security Tag Type 7

Tag Type 7, the Free Form Tag Type, is intended as a tag type that can carry a user-defined type of data appropriate for use with the protocol handling the labels. The tag usage will be domain specific but registration with IANA is recommended. Examples of data that may be conveyed with this Tag Type are human/machine readable time stamps, human-readable policy identifiers, and privacy marks.

2.3.6 Tag Type 8

This is referred to as the "Security Policy ID" tag type. The format of this tag type is as follows:

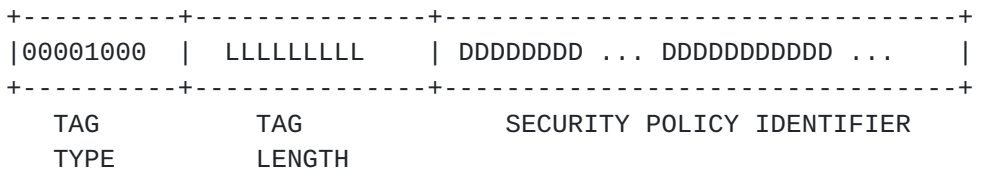

Figure 9. Tag Type 9 Format

2.3.6.1 Tag Type

This field is 1 octet in length and has a value of 8.

2.3.6.2 Tag Length

This field is 1 octet in length. It gives the total length of the tag in octets including the type and length fields.

2.3.6.3 Policy ID Field

The length of this field is variable and ranges from 4 to

16 octets and is stored in network byte order. The contents of this field is a Security Policy Identifier. The interpretation of this field is defined within a given DOI.

3. Access and Routing Control

Internet security labels are designed to enable user communities to protect information. The protection is based on access control procedures which examine the labeled information to be accessed (or forwarded) and make decisions based on the destination and its security characteristics.

The labels have been designed to make for extremely fast access control processing. They are also efficient with respect to channel bit usage. These are important factors, particularly for routers, gateways, etc. processing these labels.

3.1 Clearance Considerations

The "clearance" for a destination is the aggregate of the security categories which have been given to the destination. Access control decisions are based on the clearance of the destination and the security label attached to the file or other security object which is to be accessed.

As an example, we assume the destination has a clearance label associated in which the values in the fields are in the same format as those in the security label attached to the file to be accessed. (If not the destination label or for example, if in a given domain the first bit in a type 1 tag is assigned the restrictive clearance A and the second leftmost bit is assigned to clearance B, then the restrictive values field which gives the clearance for the destination will also have the first bit assigned to A and the second bit to B. A tag attached to information to be accessed with its bit map of categories containing 01000000 then requires a destination to have a 1 in the second position of the restrictive bit map giving the restrictive section of the clearance of the destination.

In order to make an access control decision concerning restrictive values then the access control program simply complements (inverts) the destination restrictive bit map field and ANDs the result with the bit map of categories for

the tag attached to the file. If the result is non-zero

access is refused, if the result is all 0 then access is permitted. (In this simple example, the type 1 bit map length is assumed to be 8.)

Release marking access control decisions are similarly straightforward. In both cases only two to three machine language instructions need to be executed (if the access control program is written in C a similar number of statements need to be written.) This high speed operation is generally desirable and is particularly desirable at communication levels. We note that the DOIs of the file tag and destination DOI must be the same so the correct destination tags will be selected.

3.2 Error Processing

The following table shows specific ICMP messages which have been used for error responses. These are intended for TCP/IP usage. At other layers local rules apply.

Condition	Action
ISL missing	An ICMP "parameter problem" type 12) is generated and must be returned to the originator through the same input channel from which it was received. The code field of the ICMP is set to "ISL missing" (code 1) and the ICMP pointer is set to 134.
Unrecognized label	An "ICMP parameter problem" (type 12) is generated and must be returned to the originator through the same input channel from which it was received. The ICMP code field is set to "bad parameter" (code 0) and the pointer is set to the start of the ISL field that is unrecognized.

Incoming violation	An ICMP "destination unreachable" (type 3)is generated and must be returned to the originator through the same input channel from which it was received. The

	code field of the ICMP is set to "communications with destination host administratively prohibited" (code 10).
Forwarding violation	An ICMP "destination unreachable" (type 3)is generated and returned to the originator. The code field of the ICMP is set to "communication with destination network administratively prohibited" (code 9).

3.3 Example Tag Procedures

The basic rule for processing tags is that every test for access control associated with each sensitivity tag in an ISL must be passed in order for the ISL to be forwarded. If an ISL contains a type 1 tag and a type 6 tag then the test for sensitivity hierarchical level must be passed followed by the type 1 tag bit map test followed by the release markings bit map test. Only if all tests are passed is the PDU forwarded.

The sensitivity hierarchical level test for a type 1, 2, or 5 tag consists of seeing if the binary number in the sensitivity level section is within the prescribed range. As an example, assume an ISL processing element has a stored eight bit lower range value of B(l) and a stored upper range value of B(u), where B(l) and B(u) are binary numbers with decimal values 0 to 255 and B(l) <= B(u). (Both B(l) and B(u) can be set by the system security managers.) If the ISL sensitivity level section has the value B the ISL passes only if Bl <= B <= Bu.

If the ISL carries a type 1 tag the processor will store a vector Vs (which can be set by the system security manager) which has a 1 in each position corresponding to a category the

subject can transmit, receive, or pass. Every 1 in the ISL bit map must correspond to a 1 in Vs in order for the test to succeed and the PDU to be forwarded.

If the type 2 tag is used for restrictive markings the ISL processor for a specific DOI will have a stored list of 2-octet binary values Ls = (Vs(1),Vs(2),...,Vs(m)) where each Vs(i) = (v(0),v(1),...,v(15)); v(j) = 0, 1, and each Vs(i) corresponds to a category. If we call the list of two octet

enumerated values in an ISL type 2 tag Lc = (Vc(1),Vc(2),...,Vc(m)) where the Vc(k) are 2-octet vectors, then each 2-octet vector in Lc must also be in Ls in order for the test to be passed.

For the type 5 "range" tags the same list Ls used for type 2 tags in the processor for a specific DOI is used. Then if the Top/Bottom pair Tp/Bp, where Tp and B(p) are 2-octet vectors, occurs in an ISL type 5 tag, each binary value B(j) must occur in Ls for B(p) <= B(j)<= Tp in Ls. Here the values in Ls and Tp, B(p) and B(j) are all considered to be binary values from 0 to 255.

The Release Markings Security Policy is similar to the Restrictive Security
Policy procedure.

For example, suppose we have a protocol data unit ISL with a release category list that includes just AMGEN and BIOGEN. This protocol data unit is sent to host A and host B. Host A has a release category list of NOVARTIS, ROCHE, and MERCK. The protocol data unit would be rejected since the two lists do not share at least one category. Host B, however has a list of AMGEN, ROCHE, and PATHOGENESIS. It could receive the protocol data unit because both lists share the release category AMGEN. Notice that Host B's list did not have to contain the release category BIOGEN to receive the protocol data unit.

In order for an ISL PDU to pass the access requirements it must pass the following test if it contains a type 6 tag. The access control processor will contain a binary valued vector Vr = (v(0),v(1),...,v(n)) where the v(i) = 0 or 1, associated with the DOI in the ISL. Call the release map in the ISL tag Vc=(v(0),v(1),...,v(m)), then if one or more 1s in

Vc are in the same position as 1s in Vr the test is passed. Notice the m in Vc can be less than the n in Vr since trailing octets with all 1s are not transmitted in ISL release tags. In performing tests the Vr can be truncated to the length of Vc or the ISL tag can be extended by adding 1s. Given that m = n then if the vector sum Vr + Vc contains one or more 0s the test is passed.

4. Security Considerations

This entire document discusses an encoding standard for encoding security markings.

5. References

[1] Postel, J., Internet Official Protocol Standards, STD 1, RFC 1540, Internet Architecture Board, October 1993.

[2] Atkinson, R.., Security Architecture for the Internet Protocol, RFC 1825, Cisco Systems, 10 November 1996.

[3] Atkinson, R., IP Encapsulating Security Payload, RFC-1827, 4 June 1996.

[4] Atkinson, R. IP Encapsulating Header, RFC-1826, 4 June 1996.

[5] Kent, Steve, US DoD Security Options for the Internet Protocol, RFC-1108, DDN Network Information Center, November 1991.

[6] National Institute of Standards and Technology, Standard Security Label for Information Transfer, FIPS PUB 188, November 1995.

[7] MIL-STD-2045-48501, Common Security Label (CSL), June, 1996.

Authors Addresses

Thomas C. Bartee
IDA

1801 N. Beauregard St.
Alexandria, VA 22311
Phone: (703) 845-2547
Fax: (703) 845-6722
Email: TBartee@Pop.Erols.Com

Nelson W. Alvarez
DISA/JIEO
ATTN: JEBBD
Bldg 283
Fort Monmouth, NJ 07703
Phone: (908) 427-6853
Fax: (908) 532-0853
Email: alvarezn@ftm.disa.mil

C. Dale Nunley
P.O. Box 11

Annapolis Junction
MD 20701
Phone: (301) 912-1019
Fax: (301) 912-1019
Email:dnunley@romulus.ncsc.mil