Network Working Group
Request for Comments: 5570
Category: Informational

M. StJohns
Consultant
R. Atkinson
Extreme Networks
G. Thomas
US Department of Defense

July 2009

Common Architecture Label IPv6 Security Option (CALIPSO)

Abstract

This document describes an optional method for encoding explicit packet Sensitivity Labels on IPv6 packets. It is intended for use only within Multi-Level Secure (MLS) networking environments that are both trusted and trustworthy.

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

IESG Note

This RFC specifies the use of an IPv6 hop-by-hop option. The IESG notes that general deployment of protocols with hop-by-hop options are problematic, and the development of such protocols is consequently discouraged. After careful review, the IETF has determined that a hop-by-hop option is an appropriate solution for this specific limited environment and use case. Furthermore, the mechanism specified in this RFC is only applicable to closed IP networks. It is unsuitable for use and ineffective on the global public Internet.

Table of Contents

1. Introduction

The original IPv4 specification in RFC 791 includes an option for labeling the sensitivity of IP packets. That option was revised by RFC 1038 and later by RFC 1108 [RFC791] [RFC1038] [RFC1108]. Although the IETF later deprecated RFC 1108, that IPv4 option continues to be in active use within a number of closed Multi-Level Secure (MLS) IP networks.

One or another IP Sensitivity Label option has been in limited deployment for about two decades, most usually in governmental or military internal networks. There are also some commercial sector deployments, where corporate security policies require Mandatory Access Controls be applied to sensitive data. Some banks use MLS technology to restrict sensitive information, for example information about mergers and acquisitions. This IPv6 option, like its IPv4 predecessors, is only intended for deployment within private internetworks, disconnected from the global Internet. This document specifies the explicit packet labeling extensions for IPv6 packets.

1.1. History

This document is a direct descendent of RFC 1038 and RFC 1108 and is a close cousin to the work done in the Commercial IP Security Option (CIPSO) Working Group of the Trusted Systems Interoperability Group (TSIG) [FIPS-188]. The IP Security Option defined by RFC 1038 was designed with one specific purpose in mind: to support the fielding of an IPv4 packet-encryption device called a BLACKER [RFC1038]. Because of this, the definitions and assumptions in those documents were necessarily focused on the US Department of Defense and the BLACKER device. Today, IP packet Sensitivity Labeling is most commonly deployed within Multi-Level Secure (MLS) environments, often composed of Compartmented Mode Workstations (CMWs) connected via a Local Area Network (LAN). So the mechanism defined here is accordingly more general than either RFC 1038 or RFC 1108 were.

Also, the deployment of Compartmented Mode Workstations ran into operational constraints caused by the limited, and relatively small, space available for IPv4 options. This caused one non-IETF specification for IPv4 packet labeling to have a large number of sub-options. A very unfortunate side effect of having sub-options within an IPv4 label option was that it became much more challenging to implement Intermediate System support for Mandatory Access Controls (e.g., in a router or MLS guard system) and still be able to forward traffic at, or near, wire-speed.

In the last decade or so, typical Ethernet link speeds have changed from 10 Mbps half-duplex to 1 Gbps full-duplex. The 10 Gbps full-duplex Ethernet standard is widely available today in routers, Ethernet switches, and even in some servers. The IEEE is actively developing standards for both 40 Gbps Ethernet and 100 Gbps Ethernet as of this writing. Forwarding at those speeds typically requires support from Application-Specific Integrated Circuits (ASICs); supporting more complex packet formats usually requires significantly more gates than supporting simpler packet formats. So the pressure to have a single simple option format has only increased in the past decade, and is only going to increase in the future.

When IPv6 was initially being developed, it was anticipated that the availability of IP Security, in particular the Encapsulating Security Payload (ESP) and the IP Authentication Header (AH), would obviate the need for explicit packet Sensitivity Labels with IPv6 [RFC1825] [RFC4301] [RFC4302] [RFC4303]. For MLS IPv6 deployments where the use of AH or ESP is practical, the use of AH and/or ESP is recommended.

However, some applications (e.g., distributed file systems), most often those not designed for use with Compartmented Mode Workstations or other Multi-Level Secure (MLS) computers, multiplex different transactions at different Sensitivity Levels and/or with different privileges over a single IP communications session (e.g., with the User Datagram Protocol). In order to maintain data Sensitivity Labeling for such applications, to be able to implement routing and Mandatory Access Control decisions in routers and guards on a per-IP-packet basis, and for other reasons, there is a need to have a mechanism for explicitly labeling the sensitivity information for each IPv6 packet.

Existing Layer 3 Virtual Private Network (VPN) technology can’t solve the set of issues addressed by this specification, for several independent reasons. First, in a typical deployment, many labeled packets will flow from an MLS End System through some set of networks to a receiving MLS End System. The received per-packet label is used by the receiving MLS End System to determine which Sensitivity Label to associate with the user data carried in the packet. Existing Layer 3 VPN specifications do not specify any mechanism to carry a Sensitivity Label. Second, existing Layer 3 VPN technologies are not implemented in any MLS End Systems, nor in typical single-level End System operating systems, but instead typically are only implemented in routers. Adding a Layer 3 VPN implementation to the networking stack of an MLS End System would be a great deal more work than adding this IPv6 option to that same MLS End System. Third, existing Layer 3 VPN specifications do not support the use of Sensitivity Labels to select a VPN to use in carrying a packet, which function is

essential if one wanted to obviate this IPv6 option. Substantial new standards development, along with significant new implementation work in End Systems, would be required before a Layer 3 VPN approach to these issues could be used. Developing such specifications, and then implementing them in MLS systems, would need substantially greater effort than simply implementing this IPv6 label option in an MLS End System (or in a label-aware router). Further, both the MLS user community and the MLS implementer community prefer the approach defined in this specification.

1.2. Intent and Applicability

Nothing in this document applies to any system that does not claim to implement this document.

This document describes a generic way of labeling IPv6 datagrams to reflect their particular sensitivity. Provision is made for separating data based on domain of interpretation (e.g., an agency, a country, an alliance, or a coalition), the relative sensitivity (i.e., Sensitivity Levels), and need-to-know or formal access programs (i.e., compartments or categories).

A commonly used method of encoding Releasabilities as if they were Compartments is also described. This usage does not have precisely the same semantics as some formal Releasability policies, but existing Multi-Level Secure operating systems do not contain operating system support for Releasabilities as a separate concept from compartments. The semantics for this sort of Releasability encoding is close to the formal policies and has been deployed by a number of different organizations for at least a decade now.

In particular, the authors believe that this mechanism is suitable for deployment in United Nations (UN) peace-keeping operations, in North Atlantic Treaty Organisation (NATO) or other coalition operations, in all current US Government MLS environments, and for deployment in other similar commercial or governmental environments. This option would not normally ever be visible in an IP packet on the global public Internet.

Because of the unusually severe adverse consequences (e.g., loss of life, loss of very large sums of money) likely if a packet labeled with this IPv6 Option were to escape onto the global public Internet, organizations deploying this mechanism have unusually strong incentives to configure security controls to prevent labeled packets from ever appearing on the global public Internet. Indeed, a primary purpose of this mechanism is to enable deployment of Mandatory Access Controls for IPv6 packets.

However, to ensure interoperability of both End Systems and Intermediate Systems within such a labeled deployment of IPv6, it is essential to have an open specification for this option.

This option is NOT designed to be an all-purpose label option and specifically does not include support for generic Domain Type Enforcement (DTE) mechanisms. If such a DTE label option is desired, it ought to be separately specified and have its own (i.e., different) IPv6 option number.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

1.3. Deployment Examples

Two deployment scenarios for IP packet Sensitivity Labels are most common. We should first note that in typical deployments, all people having access to an unencrypted link are cleared for all unencrypted information traversing that link. Also, MLS system administrators normally have previously been cleared to see all of the information processed or stored by that MLS system. This specification does not seek to eliminate all potential covert channels relating to this IPv6 option.

In the first scenario, all the connected nodes in a given private internetwork are trusted systems that have Multi-Level Secure (MLS) operating systems, such as Compartmented Mode Workstations (CMWs), that support per-packet Sensitivity Labels [TCSEC] [TNI] [CMW] [MLOSPP]. In this type of deployment, all IP packets carried within the private internetwork are labeled, the IP routers apply mandatory access controls (MAC) based on the packet labels and the sensitivity ranges configured into the routers, all End Systems include packet Sensitivity Labels in each originated packet, and all End Systems apply Mandatory Access Controls to each received packet. Packets received by a router or End System that have a Sensitivity Label outside the permitted range for the receiving interface (or, in the case of a router, outside the permitted range for either the incoming or the outgoing interface) are dropped because they violate the MAC policy.

The second scenario is a variation of the first, where End Systems with non-MLS operating systems are present on certain subnetworks of the private internetwork. By definition, these non-MLS End Systems operate in "system high" mode. In "system high" mode, all information on the system is considered to have the sensitivity of the most sensitive data on the system. If a system happens to contain data only at one Sensitivity Level, this would also be an

example of "system high" operation. In this scenario, each subnetwork that contains any single-level End Systems has one single "default" Sensitivity Label that applies to all single-level systems on that IP subnetwork. Because those non-MLS End Systems are unable to create packets containing Sensitivity Labels and are also unable to apply MAC enforcement on received packets, security gateways (which might, for example, be label-aware IP routers) connected to such subnetworks need to insert sensitivity labels to packets originated by the "system high" End Systems that are to be forwarded off subnet. While the CALIPSO IPv6 option is marked as "ignore if unrecognized", there are some deployed IPv6 End Systems with bugs. Users can't fix these operating system bugs; some users need to be able to integrate their existing IPv6 single-level End Systems to have a useful overall MLS deployment. So, for packets destined for IP subnetworks containing single-level End Systems, those last-hop security gateways also apply Mandatory Access Controls (MAC) and then either drop (if the packet is not permitted on that destination subnet) exclusive-or remove Sensitivity Labels and forward packets onto those "system high" subnetworks (if the packet is permitted on that destination subnetwork).

The authors are not aware of any existing MLS network deployments that use a commercial Network Address Translation (NAT), Network Address and Port Translation (NAPT), or any other commercial "middlebox" device. For example, NAT boxes aren't used, unlike practices in some segments of the public Internet.

Similarly, the authors are not aware of any existing MLS network deployments that use a commercial firewall. MLS networks normally are both physically and electronically isolated from the global Internet, so operators of MLS networks are not concerned about external penetration (e.g., by worms, viruses, or the like). Similarly, all users of the MLS network have been cleared using some process specific to that organization, and hence are believe to be trustworthy. In a typical deployment, all computers connected to the MLS network are in a physically secure room or building (e.g., protected by guards with guns). Electronic equipment that enters such a space typically does not leave. Items such as USB memory sticks are generally not permitted; in fact, often the USB ports on MLS computers have been removed or otherwise made inoperable to prevent people from adding or removing information.

Also, for security reasons, content transformation in the middle of an MLS network is widely considered undesirable, and so is not typically undertaken. Hypothetically, if such content transformation were undertaken, it would be performed by a certified MLS system that has been suitably accredited for that particular purpose in that particular deployment.

2. Definitions

This section defines several terms that are important to understanding and correctly implementing this specification. Because of historical variations in terminology in different user communities, several terms have defined synonyms.

The verb "dominate" is used in this document to describe comparison of two Sensitivity Labels within a given Domain of Interpretation. Sensitivity Label A dominates Sensitivity Label B if the Sensitivity Level of A is greater than or equal to the Sensitivity Level of B AND the Compartment Set of A is a superset (proper or improper) of the Compartment Set of B. This term has been used in Multi-Level Secure circles with this meaning for at least two decades.

2.1. Domain of Interpretation

A Domain of Interpretation (DOI) is a shorthand way of identifying the use of a particular labeling, classification, and handling system with respect to data, the computers and people who process it, and the networks that carry it. The DOI policies, combined with a particular Sensitivity Label (which is defined to have meaning within that DOI) applied to a datum or collection of data, dictates which systems, and ultimately which persons may receive that data.

In other words, a label of "SECRET" by itself is not meaningful; one also must know that the document or data belongs to some specific organization (e.g., US Department of Defense (DoD), US Department of Energy (DoE), UK Ministry of Defence (MoD), North Atlantic Treaty Organisation (NATO), United Nations (UN), a specific commercial firm) before one can decide on who is allowed to receive the data.

A CALIPSO DOI is an opaque identifier that is used as a pointer to a particular set of policies, which define the Sensitivity Levels and Compartments present within the DOI, and by inference, to the "real-world" (e.g., used on paper documents) equivalent labels (See "Sensitivity Label" below). Registering or defining a set of real-world security policies as a CALIPSO DOI results in a standard way of labeling IP data originating from End Systems "accredited" or "approved" to operate within that DOI and the constraints of those security policies. For example, if one did this for the US Department of Defense, one would list all the acceptable labels such as "SECRET" and "TOP SECRET", and one would link the CALIPSO DOI to the [DoD5200.28] and [DoD5200.1-R] documents, which define how to mark and protect data with the US Department of Defense (DoD) [DoD5200.28] [DoD5200.1-R].

The scope of the DOI is dependent on the organization creating it. In some cases, the creator of the DOI might not be identical to a given user of the DOI. For example, a multi-national organization (e.g., NATO) might create a DOI, while a given member nation or organization (e.g., UK MoD) might be using that multi-national DOI (possibly along with other DOIs created by others) within its private networks. To provide a different example, the United States might establish a DOI with specific meanings, which correspond to the normal way it labels classified documents and which would apply primarily to the US DoD, but those specific meanings might also apply to other associated agencies. A company or other organization also might establish a DOI, which applies only to itself.

NOTE WELL: A CALIPSO Domain of Interpretation is different from, and is disjoint from, an Internet Security Association and Key Management Protocol (ISAKMP) / Internet Key Exchange (IKE) Domain of Interpretation. It is important not to confuse the two different concepts, even though the terms might superficially appear to be similar.

2.2. Sensitivity Level

A Sensitivity Level represents a mandatory separation of data based on relative sensitivity. Sensitivity Levels ALWAYS have a specific ordering within a DOI. Clearance to access a specific level of data also implies access to all levels whose sensitivity is less than that level. For example, if the A, B, and C are levels, and A is more sensitive than B, which is in turn more sensitive than C (A > B > C), access to data at the B level implies access to C as well. As an example, common UK terms for a Sensitivity Level include (from low to high) "UNCLASSIFIED", "RESTRICTED", "CONFIDENTIAL", "SECRET", and "MOST SECRET".

NOTE WELL: A Sensitivity Level is only one component of a Sensitivity Label. It is important not to confuse the two terms. The term "Sensitivity Level" has the same meaning as the term "Security Level".

2.3. Compartment

A Compartment represents a mandatory segregation of data based on formal information categories, formal information compartments, or formal access programs for specific types of data. For example, a small startup company creates "FINANCE" and "R&D" compartments to protect data critical to its success -- only employees with a specific need to know (e.g., the accountants and controller for "FINANCE", specific engineers for "R&D") are given access to each compartment. Each Compartment is separate and distinct. Access to

one Compartment does not imply access to any other Compartment. Data may be protected in multiple compartments (e.g., "FINANCE" data about a new "R&D" project) at the same time, in which case access to ALL of those compartments is required to access the data. Employees only possessing clearance for a given Sensitivity Level (i.e., without having clearance for any specific compartments at that Sensitivity Level) do not have access to any data classified in any compartments (e.g., "SECRET FINANCE" dominates "SECRET").

NOTE WELL: The term "category" has the same meaning as "compartment". Some user communities have used the term "category", while other user communities have used the term "compartment", but the terms have identical meaning.

2.4. Releasability

A Releasability represents a mandatory segregation of data, based on a formal decision to release information to others.

Historically, most MLS deployments handled Releasability as if it were an inverted Compartment. Strictly speaking, this provides slightly different semantics and behavior than a paper marked with the same Releasabilities would obtain, because the formal semantics of Compartments are different from the formal semantics of Releasability. The differences in behavior are discussed in more detail later in this sub-section.

In practice, for some years now some relatively large MLS deployments have been encoding Releasabilities as if they were inverted Compartments. The results have been tolerable and those deployments are generally considered successful by their respective user communities. This description is consistent with these MLS deployments, so has significant operational experience behind it.

2.4.1. Releasability Conceptual Example

For example, two companies (ABC and XYZ) are engaging in a technical alliance. ABC labels all information present within its enterprise that is to be shared as part of the alliance as REL XYZ (e.g., COMPANY CONFIDENTIAL REL XYZ).

However, unlike the compartment example above, COMPANY CONFIDENTIAL dominates COMPANY CONFIDENTIAL REL XYZ. This means that XYZ employees granted a COMPANY CONFIDENTIAL REL XYZ clearance can only access releasable material, while ABC employees with a COMPANY CONFIDENTIAL clearance can access all information.

If REL XYZ were managed as a compartment, then users granted a COMPANY CONFIDENTIAL REL XYZ clearance would have access to all of ABC's COMPANY CONFIDENTIAL material, which is undesirable.

Releasabilities can be combined (e.g., COMPANY CONFIDENTIAL REL XYZ/ABLE). In this case, users possessing a clearance of either COMPANY CONFIDENTIAL, COMPANY CONFIDENTIAL REL XYZ, COMPANY CONFIDENTIAL REL ABLE, or COMPANY CONFIDENTIAL REL XYZ/ABLE can access this information.

2.4.2. Releasability Encoding

Individual bits in this option's Compartment Bitmap field MAY be used to encode "releasability" information. The process for making this work properly is described below.

This scheme is carefully designed so that intermediate systems need not know whether a given bit in the Compartment Bitmap field represents a compartment or a Releasability. All that an Intermediate System needs to do is apply the usual comparison (described in Section 2.5.1 and 2.5.2) to determine whether or not a packet's label is in-range for an interface. This simplifies both the configuration and implementation of a label-aware Intermediate System.

Unlike bits that represent compartments, bits that represent a Releasability are "active low".

If a given Releasability bit in the Compartment Bitmap field is "0", the information may be released to that community. If the compartment bit is "1", the information may not be released to that community.

Only administrative interfaces used to present or construct binary labels in human-readable form need to understand the distinction between Releasability bits and non-Releasability bits. Implementers are encouraged to describe Releasability encoding in the documentation supplied to users of systems that implement this specification.

2.4.2. Releasability Encoding Examples

For objects, such as IP packets, let bits 0-3 of the Compartment Bitmap field be dedicated to controlling Releasability to the communities A, B, C, and D, respectively.

Example 1: Not releasable to any community:
This is usually how handling restrictions
such as "No Foreigners (NO FORN)" are encoded.
ABCD == 1111

Example 2: Releasable only to community A and community C:
ABCD == 0101

Example 3: Releasable only to community B:
ABCD == 1011

Example 4: Releasable to communities A,B,C, & D:
ABCD == 0000

For subjects, such as clearances of users, the same bit encodings are used for Releasabilities as are used for objects (see above).

Example 1: Clearance not belonging to any community:
This user can see information belonging
to any Releasability community, since s/he
is not in any Releasability community.
ABCD = 1111

Example 2: Clearance belonging to community A and C:
This user can only see Releasable AC information,
and cannot see Releasable A information.
ABCD == 0101

Example 3: Clearance belonging to community B:
This user can only see Releasable B information.
ABCD == 1011

Example 4: Clearance belongs to communities A,B,C, and D:
This user can only see Releasable ABCD information,
and cannot (for example) see Releasable AB or
Releasable BD information.
ABCD == 0000

Now we consider example comparisons for an IP router that is enforcing MAC by using CALIPSO labels on some interface:

Let the MINIMUM label for that router interface be:
CONFIDENTIAL RELEASABLE AC

Therefore, this interface has a minimum Releasability of 0101.

Let the MAXIMUM label for that router interface be:
TOP SECRET NOT RELEASABLE

Therefore, this interface has a maximum Releasability of 1111.

For the range comparisons, the bit values for the current packet need to be "greater than or equal to" the minimum value for the interface AND also the bit values for the current packet need to be "less than or equal to" the maximum value for the interface, just as with compartment comparisons. The inverted encoding scheme outlined above ensures that the proper results occur.

Consider a packet with label CONFIDENTIAL RELEASABLE AC:

1) Sensitivity Level comparison:
   (CONFIDENTIAL <= CONFIDENTIAL <= TOP SECRET)
   so the Sensitivity Level is "within range" for that router interface.
2) Compartment bitmap comparison:
   The test is [(0101 >= 0101) AND (0101 <= 1111)],
   so the Compartment bitmap is "within range" for that router interface.

Consider a packet with label CONFIDENTIAL RELEASABLE ABCD:

1) Sensitivity Label comparison:
   (CONFIDENTIAL <= CONFIDENTIAL <= TOP SECRET)
   so the Sensitivity Level is "within range" for that router interface.
2) Compartment bitmap comparison:
   The test is [(0000 >= 0101) AND (0000 <= 1111)],
   so the Compartment Bitmap is NOT "within range" for that router interface.

Consider a packet with label SECRET NOT RELEASABLE:

1) Sensitivity Label comparison:
   (CONFIDENTIAL <= SECRET <= TOP SECRET)
   so the Sensitivity Level is "within range" for that router interface.
2) Compartment bitmap comparison:
   The test is [(1111 >= 0101) AND (1111 <= 1111)],
   so the Compartment bitmap is "within range" for that router interface.

2.4.3. Limitations of This Releasability Approach

For example, if one considers a person "Jane Doe" who is a member of two Releasability communities (A and also B), she is permitted to see a paper document that is marked "Releasable A", "Releasable B", or "Releasable AB" -- provided that her Clearance and Compartments are in-range for the Sensitivity Level and Compartments (respectively) of the paper document.

Now, let us consider an equivalent electronic example implemented and deployed as outlined above. In this, we consider two Releasability communities (A and B). Those bits will be set to 00 for the electronic user ID used by user "Jane Doe".

However, the electronic Releasability approach above will ONLY permit her to see information marked as "Releasable AB". The above electronic approach will deny her the ability to read documents marked "Releasable A" or "Releasable B". This is because "Releasable A" is encoded as "01", "Releasable B" is encoded as "10", while "Releasable AB" is encoded as "00". If one looks at the compartment dominance computation, "00" dominates "00", but "00" does NOT dominate "01", and "00" also does NOT dominate "10".

Users report that the current situation is tolerable, but not ideal. Users also report that various operational complexities can arise from this approach.

Several deployments work around this limitation by assigning an electronic user several parallel clearances. Referring to the (fictitious) example above, the user "Jane Doe" might have one clearance without any Releasability, another separate clearance with Releasability A, and a third separate clearance with Releasability B. While this has implications (e.g., a need to be able to associate multiple separate parallel clearances with a single user ID) for implementers of MLS systems, this specification cannot (and does not) levy any requirements that an implementation be able to associate multiple clearances with each given user ID because that level of detail is beyond the scope of an IP labeling option.

Separating the Releasability bits into a separate bitmap within the CALIPSO option was seriously considered. However, existing MLS implementations lack operating system support for Releasability. So even if CALIPSO had a separate bitmap field, those bits would have been mapped to Compartment bits by the sending/receiving nodes, so the operational results would not have been different than those described here.

Several MLS network deployments connect MLS End Systems both to a labeled national network and also to a labeled coalition network simultaneously. Depending on whether the data is labeled according to national rules or according to coalition rules, the set of Releasability marks will vary. Some choices are likely to lead to more (or fewer) incorrect Releasability decisions (although the results of the above Releasability encodings are believed to be fail-safe).

2.5. Sensitivity Label

A Sensitivity Label is a quadruple consisting of a DOI, a Sensitivity Level, a Compartment Set, and a Releasability Set. The Compartment Set may be the empty set if and only if no compartments apply. A Releasability Set may be the empty set if and only if no Releasabilities apply. A DOI used within an End System may be implicit or explicit depending on its use. CALIPSO Sensitivity Labels always have an explicit DOI. A CALIPSO Sensitivity Label consists of a Sensitivity Label in a particular format (defined below). A CALIPSO Sensitivity Label ALWAYS contains an explicit DOI value. In a CALIPSO Sensitivity Label, the Compartment Bitmap field is used to encode both the logical Compartment Set and also the logical Releasability Set.

End Systems using operating systems with MLS capabilities that also implement IPv6 normally will be able to include CALIPSO labels in packets they originate and will be able to enforce MAC policy on the CALIPSO labels in any packets they receive.

End Systems using an operating system that lacks Multi-Level Secure capabilities operate in "system high" mode. This means that all data on the system is considered to have the Sensitivity Label of the most sensitive data on the system. Such a system normally is neither capable of including CALIPSO labels in packets that it originates nor of enforcing CALIPSO labels in packets that it receives.

NOTE WELL: The term "Security Marking" has the same meaning as "Sensitivity Label".

2.5.1. Sensitivity Label Comparison

Two Sensitivity Labels (A and B) can be compared. Indeed, Sensitivity Labels exist primarily so they can be compared as part of a Mandatory Access Control decision. Comparison is critical to determining if a subject (a person, network, etc.) operating at one Sensitivity Label (A) should be allowed to access an object (file, packet, route, etc.) classified at another Sensitivity Label (B). The comparison of two labels (A and B) can return one (and only one) of the following results:

1) A dominates B (e.g., A=SECRET, B=UNCLASSIFIED);
   A can read B,

2) B dominates A (e.g., A=UNCLASSIFIED, B=SECRET);
   A cannot access B,

3) A equals B (e.g., A=SECRET, B=SECRET);
A can read/write B,

exclusive-or

4) A is incomparable to B (e.g., A=SECRET R&D, B=SECRET FINANCE);
A cannot access B, and also, B cannot access A.

By definition, if A and B are members of different DOIs, the result of comparison is always incomparable. It is possible to overcome this if and only if A and/or B can be translated into some common DOI, such that the labels are then interpretable.

2.5.2. Sensitivity Label Range

A range is a pair of Sensitivity Labels, which indicate both a minimum and a maximum acceptable Sensitivity Label for objects compared against it. A range is usually expressed as "<minimum> : <maximum>" and always has the property that the maximum Sensitivity Label dominates the minimum Sensitivity Label. In turn, this requires that the two Sensitivity Labels MUST be comparable.

A range where <minimum> equals <maximum> may be expressed simply as "<minimum>"; in this case, the only acceptable Sensitivity Label is <minimum>.

2.6. Import

The act of receiving a datagram and translating the CALIPSO Sensitivity Label of that packet into the appropriate internal (i.e., end-system-specific) Sensitivity Label.

2.7. Export

The act of selecting an appropriate DOI for an outbound datagram, translating the internal (end-system-specific) label into an CALIPSO Sensitivity Label based on that DOI, and sending the datagram. The selection of the appropriate DOI may be based on many factors including, but not necessarily limited to:

Source Port
Destination Port
Transport Protocol
Application Protocol
Application Information
End System
Subnetwork
Network
Sending Interface
System Implicit/Default DOI

Regardless of the DOI selected, the Sensitivity Label of the outbound datagram must be consistent with the security policy monitor of the originating system and also with the DOI definition used by all other devices cognizant of that DOI.

2.8. End System

An End System is a host or router from which a datagram originates or to which a datagram is ultimately delivered.

The IPv6 community has defined the term Node to include both Intermediate Systems and End Systems [RFC2460].

2.9. Intermediate System

An Intermediate System (IS) is a node that receives and transmits a particular datagram without being either the source or destination of that datagram. An Intermediate System might also be called a "gateway", "guard", or "router" in some user communities.

So an IPv6 router is one example of an Intermediate System. A firewall or security guard device that applies security policies and forwards IPv6 packets that comply with those security policies is another example of an Intermediate System.

An Intermediate System may handle ("forward") a datagram destined for some other node without necessarily importing or exporting the datagram to/from itself.

NOTE WELL: Any given system can be both an End System and an Intermediate System -- which role the system assumes at any given time depends on the address(es) of the datagram being considered and the address(es) associated with that system.

2.10. System Security Policy

A System Security Policy (SSP) consists of a Sensitivity Label and the organizational security policies associated with content labeled with a given security policy. The SSP acts as a bridge between how the organization's Mandatory Access Control (MAC) policy is stated and managed and how the network implements that policy. Typically, the SSP is a document created by the Information Systems Security Officer (ISSO) of the site or organization covered by that SSP.

3. Architecture

This document describes a convention for labeling an IPv6 datagram within a particular system security policy. The labels are designed for use within a Mandatory Access Control (MAC) system. A real-world example is the security classification system in use within the UK Government. Some data held by the government is "classified", and is therefore restricted by law to those people who have the appropriate "clearances".

Commercial examples of information labeling schemes also exist [CW87]. For example, one global electrical equipment company has a formal security policy that defines six different Sensitivity Levels for its internal data, ranging from "Class 1" to "Class 6" information. Some financial institutions use multiple compartments to restrict access to certain information (e.g., "mergers and acquisitions", "trading") to those working directly on those projects and to deny access to other groups within the company (e.g., equity trading). A CALIPSO Sensitivity Label is the network instantiation of a particular information security policy, and the policy's related labels, classifications, compartments, and Releasabilities.

Some years ago, the Mandatory Access Control (MAC) policy for US Government classified information was specified formally in mathematical notation [BL73]. As it happens, many other organizations or governments have the same basic Mandatory Access Control (MAC) policy for information with differing ("vertical") Sensitivity Levels. This document builds upon the formal definitions of Bell-LaPadula [BL73]. There are two basic principles: "no write down" and "no read up".

The first rule means that an entity having minimum Sensitivity Level X must not be able to write information that is marked with a Sensitivity Level below X. The second rule means that an entity having maximum Sensitivity Level X must not be able to read information having a Sensitivity Level above X. In a normal deployment, information downgrading ("write down") must not occur automatically, and is permitted if and only if a person with

appropriate "downgrade" privilege manually verifies the information is permitted to be downgraded before s/he manually relabels (i.e., "downgrades") the information. Subsequent to the original work by Bell and LaPadula in this area, this formal model was extended to also support ("horizontal") Compartments of information.

This document extends Bell-LaPadula to accommodate the notion of separate Domains of Interpretation (DOI) [BL73]. Each DOI constitutes a single comparable domain of Sensitivity Labels as stated by Bell-LaPadula. Sensitivity Labels from different domains cannot be directly compared using Bell-LaPadula semantics.

This document is focused on providing specifications for (1) encoding Sensitivity Labels in packets, and (2) how such Sensitivity Labels are to be interpreted and enforced at the IP layer. This document recognizes that there are several kinds of application processing that occur above the IP layer that significantly impact end-to-end system security policy enforcement, but are out of scope for this document. In particular, how the network labeling policy is enforced within processing in an End System is critical, but is beyond the scope of a network (IP) layer Sensitivity Label encoding standard. Other specifications exist, which discuss such details [TCSEC] [TNI] [CMW] [ISO-15408] [CC] [MLOSPP].

This specification does not preclude an End System capable of providing labeled packets across some range of Sensitivity Labels. A Compartmented Mode Workstation (CMW) is an example of such an End System [CMW]. This is useful if the End System is capable of, and accredited to, separate processing across some range of Sensitivity Labels. Such a node would have a range associated with it within the network interface connecting the node to the network. As an example, an End System has the range "SECRET: TOP SECRET" associated with it in the Intermediate System to which the node is attached. SECRET processing on the node is allowed to traverse the network to other "SECRET : SECRET" segments of the network, ultimately to a "SECRET : SECRET" node. Likewise, TOP SECRET processing on the node is allowed to traverse a network through "TOP SECRET: TOP SECRET" segments, ultimately to some "TOP SECRET: TOP SECRET" node. The node in this case can allow a user on this node to access SECRET and TOP SECRET resources, provided the user holds the appropriate clearances and has been correctly configured.

With respect to a given network, each distinct Sensitivity Label represents a separate virtual network, which shares the same physical network. There are rules for moving information between the various virtual networks. The model we use within this document is based on

the Bell-LaPadula model, but is extended to cover the concept of differing Domains of Interpretation. Nodes that implement this protocol MUST enforce this mandatory separation of data.

CALIPSO provides for both horizontal ("Compartment") and vertical ("Sensitivity Level") separation of information, as well as separation based on DOI. The basic rule is that data MUST NOT be delivered to a user or system that is not approved to receive it.

NOTE WELL: Wherever we say "not approved", we also mean "not cleared", "not certified", and/or "not accredited" as applicable in one's operational community.

This specification does not enable AUTOMATIC relabeling of information, within a DOI or to a different DOI. That is, neither automatic "upgrading" nor automatic "downgrading" of information are enabled by this specification. Local security policies might allow some limited downgrading, but this normally requires the intervention of some human entity and is usually done within an End System with respect to the internal Sensitivity Label, rather than on a network or in an intermediate-system (e.g., router, guard). Automatic downgrading is not suggested operational practice; further discussion of downgrading is outside the scope of this protocol specification.

Implementers of this specification MUST NOT permit automatic upgrading or downgrading of information in the default configuration of their implementation. Implementers MAY add a configuration knob that would permit a System Security Officer holding appropriate privilege to enable automatic upgrading or downgrading of information. If an implementation supports such a knob, the existence of the configuration knob must be clearly documented and the default knob setting MUST be that automatic upgrading or downgrading is DISABLED. Automatic information upgrading and downgrading is not recommended operational practice.

Many existing MLS deployments already use (and operationally need to use) more than one DOI concurrently. User feedback from early versions of this specification indicates that it is common at present for a single network link (i.e., IP subnetwork) to carry traffic for both a particular coalition (or joint-venture) activity and also for the government (or other organization) that owns and operates that particular network link. On such a link, one CALIPSO DOI would typically be used for the coalition traffic and some different CALIPSO DOI would typically be used for non-coalition traffic (i.e., traffic that is specific to the government that owns and operates that particular network link). For example, a UK military network that is part of a NATO deployment might have and use a UK MoD DOI for

information originating/terminating on another UK system, while concurrently using a different NATO DOI for information originating/terminating on a non-UK NATO system.

Additionally, operational experience with existing MLS systems has shown that if a system only supports a single DOI at a given time, then it is impossible for a deployment to migrate from using one DOI value to a different DOI value in a smooth, lossless, zero downtime, manner.

Therefore, a node that implements this specification MUST be able to support at least two CALIPSO DOIs concurrently. Support for more than two concurrent CALIPSO DOIs is encouraged. This requirement to support at least two CALIPSO DOIs concurrently is not necessarily an implementation constraint upon MLS operating system internals that are unrelated to the network.

Indeed, use of multiple DOIs is also operationally useful in deployments having a single administration that also have very large numbers of compartments. For example, such a deployment might have one set of related compartments in one CALIPSO DOI and a different set of compartments in a different CALIPSO DOI. Some compartments might be present in both DOIs, possibly at different bit positions of the compartment bitmap in different DOIs. While this might make some implementations more complex, it might also be used to reduce the typical size of the IPv6 CALIPSO option in data packets.

Moving information between any two DOIs is permitted -- if and only if -- the owners of the DOIs:

1) Agree to the exchange,

AND

2) Publish a document with a table of equivalencies that maps the CALIPSO values of one DOI into the other and make that document available to security administrators of MLS systems within the deployment scope of those two DOIs.

The owners of two DOIs may choose to permit the exchange on or between any of their systems, or may restrict exchange to a small subset of the systems they own/accredit. One-way agreements are permissible, as are agreements that are a subset of the full table of equivalences. Actual administration of inter-DOI agreements is outside the scope of this document.

When data leaves an End System it is exported to the network, and marked with a particular DOI, Sensitivity Level, and Compartment Set. (This triple is collectively termed a Sensitivity Label.) This Sensitivity Label is derived from the internal Sensitivity Label (the end-system-specific implementation of a given Sensitivity Label), and the Export DOI. Selection of the Export DOI is described in detail in Section 6.2.1.

When data arrives at an End System, it is imported from the network to the End System. The data from the datagram takes on an internal Sensitivity Label based on the Sensitivity Label contained in the datagram. This assumes the datagram is marked with a recognizable DOI, there is a corresponding internal Sensitivity Label equivalent to the CALIPSO Sensitivity Label, and the datagram is "within range" for the receiving logical interface.

A node has one or more physical interfaces. Each physical interface is associated with a physical network segment used to connect the node, router, LAN, or WAN. One or more Sensitivity Label ranges are associated with each physical network interface. Sensitivity Label ranges from multiple DOIs must be enumerated separately. Multiple ranges from the same DOI are permissible.

Each node also might have one or more logical network interfaces.

A given logical network interface might be associated with more than one physical interface. For example, a switch/router might have two separate Ethernet ports that are associated with the same Virtual Local Area Network (VLAN), where that one VLAN mapped to a single IPv6 subnetwork [IEEE802.1Q].

A given physical network interface might have more than one associated logical interface. For example, a node might have 2 logical network interfaces, each for a different IP subnetwork ("super-netting"), on a single physical network interface (e.g., on a single Network Interface Card of a personal computer). Alternatively, also as an example, a single Ethernet port might have multiple Virtual LANs (VLANs) associated with it, where each VLAN could be a separate logical network interface.

One or more Sensitivity Label ranges are associated with each logical network interface. Sensitivity Label ranges from multiple DOIs must be enumerated separately. Multiple ranges from the same DOI are permissible. Each range associated with a logical interface must fall within a range separately defined for the corresponding physical interface.

There is specific user interest in having IPv6 routers that can apply per-logical-interface mandatory access controls based on the contents of the CALIPSO Sensitivity Labels in IPv6 packets. The authors note that since the early 1990s, and continuing through today, some commercial IPv4 router products provide MAC enforcement for the RFC 1108 IP Security Option.

In transit, a datagram is handled based on its CALIPSO Sensitivity Label, and is usually neither imported to or exported from the various Intermediate Systems it transits. There also is the concept of "CALIPSO Gateways", which import data from one DOI and export it to another DOI such that the effective Sensitivity Label is NOT changed, but is merely represented using a different DOI. In other words, such devices would be trustworthy, trusted, and authorized to provide on-the-fly relabeling of packets at the boundaries between complete systems of End Systems within a single DOI. Typically, such systems require specific certification(s) and accreditation(s) before deployment or use.

4. Defaults

This Section describes the default behavior of CALIPSO-compliant End Systems and Intermediate Systems. Implementers MAY implement configuration knobs to vary from this behavior, provided that the default behavior (i.e., if the system administrator does not explicitly change the configured behavior of the device) is as described below. If implementers choose to implement such configuration knobs, the configuration parameters and the behaviors that they enable and disable SHOULD be documented for the benefit of system administrators of those devices.

Each Intermediate System or End System is responsible for properly interpreting and enforcing the MLS Mandatory Access Control policy. Practically, this means that each node must evaluate the label on the inbound packet, ensure that this Sensitivity Label is valid (i.e., within range) for the receiving interface, and at a minimum only forward the packet to an interface and node where the Sensitivity Label of the packet falls within the assigned range of that node’s receiving interface.

Packets with an invalid (e.g., out-of-range) Sensitivity Label for the receiving interface MUST be dropped upon receipt. A Sensitivity Label is valid if and only if the Sensitivity Label falls within the range assigned to the transmitting interface on the sending system and within the range assigned to the receiving interface on the receiving system. These rules also need to be applied by Intermediate Systems on each hop that a CALIPSO-labeled packet traverses, not merely at the end points of a labeled IP session. As

an example, it is a violation of the default MLS MAC policy for a packet with a higher Sensitivity Level (e.g., "MOST SECRET") to transit a link whose maximum Sensitivity Level is less than that first Sensitivity Level (e.g., "SECRET").

If an unlabeled packet is received from a node that does not support CALIPSO Sensitivity Labels (i.e., unable to assign Sensitivity Labels itself) and the packet is destined for a node that supports CALIPSO Sensitivity Labels, then the receiving intermediate system needs to insert a Sensitivity Label. This Sensitivity Label MUST be equal to the maximum Sensitivity Label assigned to the originating node if and only if that is known to the receiving node. If this receiving Intermediate System does not know which Sensitivity Label is assigned to the originating node, then the maximum Sensitivity Label of the interface that received the unlabeled packet MUST be inserted.

NOTE WELL: The procedure in the preceding paragraph is NOT a label upgrade -- because it is not changing an existing label; instead, it is simply inserting a Sensitivity Label that has the only "safe" value, given that no other information is known to the receiving node. In large-scale deployments, it is very unlikely that a given node will have any authoritative a priori information about the security configuration of any node that is NOT on a directly attached link.

If a packet is to be sent to a node that is defined to not be Sensitivity Label aware, from a node that is label aware, then the Sensitivity Label MAY be removed upon transmission if and only if local security policy explicitly permits this. The originating node is still responsible for ensuring that the Sensitivity Label on the packet falls within the Sensitivity Label range associated with the receiving node. If the packet will traverse more than one subnetwork between origin and destination, and those subnetworks are labeled, then the packet SHOULD normally contain a Sensitivity Label so that the packet will be able to reach the destination and the Intermediate Systems will be able to apply the requisite MAC policy to the packet.

NOTE WELL: In some IPv4 MLS network deployments that exist as of the publication date, if a first-hop router receives an unlabeled IPv4 packet, the router inserts an appropriate Sensitivity Label into that IPv4 packet, in the manner described above. So sending a packet without a label across a multiple subnetwork path to a destination does not guarantee that the packet will arrive containing no Sensitivity Label.

5. Format

This section describes the format of the CALIPSO option for use with IPv6 datagrams. CALIPSO is an IPv6 Hop-By-Hop Option, rather than an IPv6 Destination Option, to ensure that a security gateway or router can apply access controls to IPv6 packets based on the CALIPSO label carried by the packet.

An IPv6 datagram that has not been tunneled contains at most one CALIPSO label. In the special case where (1) a labeled IPv6 datagram is tunneled inside another labeled IPv6 datagram AND (2) IP Security is NOT providing confidentiality protection for the inner packet, the outer CALIPSO Sensitivity Label must have the same meaning as the inner CALIPSO Sensitivity Label. For example, it would be invalid to encapsulate an unencrypted IPv6 packet with a Sensitivity Label of (SECRET, no compartments) inside a packet with an outer Sensitivity Label of (UNCLASSIFIED).

If the inner IPv6 packet is tunneled inside the Encapsulating Security Payload (ESP) and confidentiality is being provided to that inner packet, then the outer packet MAY have a different CALIPSO Sensitivity Label -- subject to local security policy.

As a general principle, the meaning of the Sensitivity Labels must be identical when one has a labeled cleartext IP packet that has been encapsulated (tunneled) inside another labeled IP packet. This is true whether one has IPv6 tunneled in IPv6, IPv4 tunneled in IPv6, or IPv6 tunneled in IPv4. This is essential to maintaining proper Mandatory Access Controls.

This option’s syntax has been designed with intermediate systems in mind. It is now common for an MLS network deployment to contain an Intermediate Systems acting as a guard (sometimes several acting as guards). Such a guard device needs to be able to very rapidly parse the Sensitivity Label in each packet, apply ingress interface MAC policy, forward the packet while aware of the packet’s Sensitivity Label, and then apply egress interface MAC policy.

At least one prior IP Sensitivity Label option [FIPS-188] used a syntax that was unduly complex to parse in IP routers, hence that option never was implemented in an IP router. So there is a deliberate effort here to choose a streamlined option syntax that is easy to parse, encode, and implement in more general terms.

5.1. Option Format

The CALIPSO option is an IPv6 Hop-by-Hop Option and is designed to comply with IPv6 optional header rules. Following the nomenclature of Section 4.2 of RFC 2460, the Option Type field of this option must have 4n+2 alignment [RFC2460].

The CALIPSO Option Data MUST NOT change en route, except when (1) "DOI translation" is performed by a trusted Intermediate System, (2) a CALIPSO Option is inserted by a trusted Intermediate System upon receipt of an unlabeled IPv6 packet, or (3) a CALIPSO Option is removed by a last-hop trusted Intermediate System immediately prior to forwarding the packet to a destination node that does not implement support for CALIPSO labels. The details of these three exceptions are described elsewhere in this document.

If the option type is not recognized by a node examining the packet, the option is ignored. However, all implementations of this specification MUST be able to recognize this option and therefore MUST NOT ignore this option if it is present in an IPv6 packet.

This option is designed to comply with the IPv6 optional header rules [RFC2460]. The CALIPSO option is always carried in a Hop-By-Hop Option Header, never in any other part of an IPv6 packet. This rule exists because IPv6 routers need to be able to see the CALIPSO label so that those routers are able to apply MLS Mandatory Access Controls to those packets.

The diagram below shows the CALIPSO option along with the required (first) two fields of the Hop-By-Hop Option Header that envelops the CALIPSO option. The design of the CALIPSO option is arranged to avoid the need for 16 bits of padding between the HDR EXT LEN field and the start of the CALIPSO option. Also, the CALIPSO Domain of Interpretation field is laid out so that it normally will be 32-bit aligned.

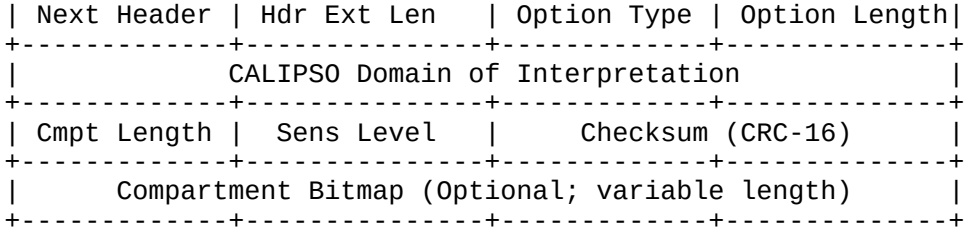

5.1.1. Option Type Field

This field contains an unsigned 8-bit value. Its value is 00000111 (binary).

Nodes that do not recognize this option should ignore it. In many cases, not all routers in a given MLS deployment will contain support for this CALIPSO option. For interoperability reasons, it is important that routers that do not support the CALIPSO forward this packet normally, even though those routers do not recognize the CALIPSO option.

In the event the IPv6 packet is fragmented, this option MUST be copied on fragmentation. Virtually all users want the choice of using the IP Authentication Header (a) to authenticate this option and (b) to bind this option to the associated IPv6 packet.

5.1.2. Option Length Field

This field contains an unsigned integer one octet in size. Its minimum value is eight (e.g., when the Compartment Bitmap field is absent). This field specifies the Length of the option data field of this option in octets. The Option Type and Option Length fields are not included in the length calculation.

5.1.3. Compartment Length Field

This field contains an unsigned 8-bit integer. The field specifies the size of the Compartment Bitmap field in 32-bit words. The minimum value is zero, which is used only when the information in this packet is not in any compartment. (In that situation, the CALIPSO Sensitivity Label has no need for a Compartment Bitmap). Note that measuring the Compartment Bitmap field length in 32-bit words permits the header to be 64-bit aligned, following IPv6 guidelines, without wasting 32 bits. Using 64-bit words for the size of the Compartment Bitmap field length would force 32 bits of padding with every option in order to maintain 64-bit alignment; wasting those bits in every CALIPSO option is undesirable.

Because this specification represents Releasabilities on the wire as inverted Compartments, the size of the Compartment Bitmap field needs to be large enough to hold not only the set of logical Compartments, but instead to hold both the set of logical Compartments and the set of logical Releasabilities.

Recall that the overall length of this option MUST follow IPv6 optional header rules, including the word alignment rules. This has implications for the valid values for this field. In some cases, the

length of the Compartment Bitmap field might need to exceed the number of bits required to hold the sum of the logical Compartments and the logical Releasabilities, in order to comply with IPv6 alignment rules.

5.1.5. Domain of Interpretation Field

This field contains an unsigned 32-bit integer. IANA maintains a registry with assignments of the DOI values used in this field. The DOI identifies the rules under which this datagram must be handled and protected. The NULL DOI, in which this field is all zeros, MUST NOT appear in any IPv6 packet on any network.

NOTE WELL: The Domain Of Interpretation value where all 4 octets contain zero is defined to be the NULL DOI. The NULL DOI has no compartments and has a single level whose value and CALIPSO representation are each zero. The NULL DOI MUST NOT ever appear on the wire. If a packet is received containing the NULL DOI, that packet MUST be dropped and the event SHOULD be logged as a security fault.

5.1.6. Sensitivity Level Field

This contains an unsigned 8-bit value. This field contains an opaque octet whose value indicates the relative sensitivity of the data contained in this datagram in the context of the indicated DOI. The values of this field MUST be ordered, with 00000000 being the lowest Sensitivity Level and 11111111 being the highest Sensitivity Level.

However, in a typical deployment, not all 256 Sensitivity Levels will be in use. So the set of valid Sensitivity Level values depends upon the CALIPSO DOI in use. This sensitivity ordering rule is necessary so that Intermediate Systems (e.g., routers or MLS guards) will be able to apply MAC policy with minimal per-packet computation and minimal configuration.

5.1.7. 16-Bit Checksum Field

This 16-bit field contains the a CRC-16 checksum as defined in Appendix C of RFC 1662 [RFC1662]. The checksum is calculated over the entire CALIPSO option in this packet, including option header, zeroed-out checksum field, option contents, and any required padding zero bits.

The checksum MUST always be computed on transmission and MUST always be verified on reception. This checksum only provides protection against accidental corruption of the CALIPSO option in cases where

neither the underlying medium nor other mechanisms, such as the IP Authentication Header (AH), are available to protect the integrity of this option.

Note that the checksum field is always required, even when other integrity protection mechanisms (e.g., AH) are used. This method is chosen for its reliability and simplicity in both hardware and software implementations, and because many implementations already support this checksum due to its existing use in various IETF specifications.

5.1.8. Compartment Bitmap Field

This contains a variable number of 64-bit words. Each bit represents one compartment within the DOI. Each "1" bit within an octet in the Compartment Bitmap field represents a separate compartment under whose rules the data in this packet must be protected. Hence, each "0" bit indicates that the compartment corresponding with that bit is not applicable to the data in this packet. The assignment of identity to individual bits within a Compartment Bitmap for a given DOI is left to the owner of that DOI.

This specification represents a Releasability on the wire as if it were an inverted Compartment. So the Compartment Bitmap holds the sum of both logical Releasabilities and also logical Compartments for a given DOI value. The encoding of the Releasabilities in this field is described elsewhere in this document. The Releasability encoding is designed to permit the Compartment Bitmap evaluation to occur without the evaluator necessarily knowing the human semantic associated with each bit in the Compartment Bitmap. In turn, this facilitates the implementation and configuration of Mandatory Access Controls based on the Compartment Bitmap within IPv6 routers or guard devices.

5.2. Packet Word Alignment Considerations

The basic option is variable length, due to the variable length Compartment Bitmap field.

Intermediate Systems that lack custom silicon processing capabilities and most End Systems perform best when processing fixed-length, fixed-location items. So the IPv6 base specification levies certain requirements on all IPv6 optional headers.

The CALIPSO option must maintain this IPv6 64-bit alignment rule for the option overall. Please note that the Compartment Bitmap field has a length in quanta of 32-bit words (e.g., 0 bits, 32 bits, 64 bits, 96 bits), which permits the overall CALIPSO option length to be 64-bit aligned -- without requiring 32 bits of NULL padding with every CALIPSO option.

6. Usage

This section describes specific protocol processing steps required for systems that claim to implement or conform with this specification.

6.1. Sensitivity Label Comparisons

This section describes how comparisons are made between two Sensitivity Labels. Implementing this comparison correctly is critical to the MLS system providing the intended Mandatory Access Controls (MACs) to network traffic entering or leaving the system.

A Sensitivity Label consists of a DOI, a Sensitivity Level, and zero or more Compartments. The following notation will be used:

```
A.DOI  = the DOI portion of Sensitivity Label A
A.LEV  = the Sensitivity Level portion of Sensitivity Label A
A.COMP = the Compartments portion of Sensitivity Label A
```

6.1.1. "Within Range"

A Sensitivity Label "M" is "within range" for a particular range "LO:HI" if and only if:

1. M, LO, and HI are members of the same DOI.

   (M.DOI == LO.DOI == HI.DOI)

2. The range is a valid range. A given range LO:HI is valid if and only if HI dominates LO.

   ((LO.LEV <= HI.LEV) && (LO.COMP <= HI.COMP))

3. The Sensitivity Level of M dominates the low-end (LO) Sensitivity Level AND the Sensitivity Level of M is dominated by the high-end (HI) Sensitivity Level.

   (LO.LEV <= M.LEV <= HI.LEV)

AND

4. The Sensitivity Label M has a Compartment Set that dominates the Compartment Set contained in the Sensitivity Label from the low-end range (LO), and that is dominated by the Compartment Set contained in the high-end Sensitivity Label (HI) from the range.

   (LO.COMP <= M.COMP <= HI.COMP)

6.1.2. "Less Than" or "Below Range"

A Sensitivity Label "M" is "less than" some other Sensitivity Label "LO" if and only if:

1. The DOI for the Sensitivity Label M is identical to the DOI for both the low-end and high-end of the range.

   (M.DOI == LO.DOI == HI.DOI)

AND EITHER

2. The Sensitivity Level of M is less than the Sensitivity Level of LO.

   (M.LEV < LO.LEV)

OR

3. The Compartment Set of Sensitivity Label M is dominated by the Compartment Set of Sensitivity Label LO.

   (M.COMP <= LO.COMP)

A Sensitivity Label "M" is "below range" for a Sensitivity Label "LO:HI", if LO dominates M and LO is not equal to M.

6.1.3. "Greater Than" or "Above Range"

A Sensitivity Label "M" is "greater than" some Sensitivity Label "HI" if and only if:

1. Their DOI’s are identical.

   (M.DOI == HI.DOI)

AND EITHER

2A. M's Sensitivity Level is above HI's Sensitivity Level.

(M.LEV > HI.LEV)

OR

2B. M's Compartment Set is greater than HI's Compartment Set.

(M.COMP > HI.COMP)

A Sensitivity Label "M" is "above range" for a Sensitivity Label, "LO:HI", if M dominates HI and M is not equal to HI.

6.1.4. "Equal To"

A Sensitivity Label "A" is "equal to" another Sensitivity Label "B" if and only if:

1. They have the exact same DOI.

   (A.DOI == B.DOI)

2. They have identical Sensitivity Levels.

   (A.LEV == B.LEV)

3. Their Compartment Sets are identical.

   (A.COMP == B.COMP)

6.1.5. "Disjoint" or "Incomparable"

A Sensitivity Label "A" is disjoint from another Sensitivity Label "B" if any of these conditions are true:

1. Their DOI's differ.

   (A.DOI <> B.DOI)

2. B does not dominate A, A does not dominate B, and A is not equal to B.

   (^((A < B) || (A > B) || (A == B)))

3. Their Compartment Sets are disjoint from each other; A's Compartment Set does not dominate B's Compartment Set AND B's Compartment Set does not dominate A's Compartment Set.

(^((A.COMP >= B.COMP) || (A.COMP <= B.COMP)))

6.2. End System Processing

This section describes CALIPSO-related processing for IPv6 packets imported or exported from an End System claiming to implement or conform with this specification. This document places no additional requirements on IPv6 nodes that do not claim to implement or conform with this document.

6.2.1. Export

An End System that sends data to the network is said to "export" it to the network. Before a datagram can leave an end system and be transmitted over a network, the following ordered steps must occur:

1. Selection of the export DOI:

   a) If the upper-level protocol selects a DOI, then that DOI is selected.

   b) Else, if there are tables defining a specific default DOI for the specific destination End System address or for the network address, then that DOI is selected.

   c) Else, if there is a specific DOI associated with the sending logical interface (i.e., IP address), then that DOI is selected.

   d) Else the default DOI for the system is selected.

NOTE WELL: A connection-oriented transport-layer protocol session (e.g., Transmission Control Protocol (TCP) session, Stream Control Transmission Protocol (SCTP) session) MUST have the same DOI and same Sensitivity Label for the life of that connection. The DOI is selected at connection initiation and MUST NOT change during the session.

A trusted multi-level application that possesses appropriate privilege MAY use multiple connection-oriented transport-layer protocol sessions with differing Sensitivity Labels concurrently.

Some trusted UDP-based applications (e.g., remote procedure call service) multiplex different transactions having different Sensitivity Levels in different packets for the same IP session (e.g., IP addresses and UDP ports are constant for a given UDP session). In such cases, the Trusted Computing Base MUST ensure that each packet is labeled with the correct Sensitivity Label for the information carried in that particular packet.

In the event the End System selects and uses a specific DOI and that DOI is not recognized by the originating node's first-hop router, the packet MUST be dropped by the first-hop router. In such a case, the networking API should indicate the connection failure (e.g., with some appropriate error, such as ENOTREACH). This fault represents (1) incorrect configuration of either the Intermediate System or of the End System or (2) correct operation for a node that is not permitted to send IPv6 packets with that DOI through that Intermediate System.

When an MLS End System is connected to an MLS LAN, it is possible that there would be more than one first-hop Intermediate System concurrently, with different Intermediate Systems having different valid Sensitivity Label ranges. Thoughtful use of the IEEE 802 Virtual LAN (VLAN) standard (e.g., with different VLAN IDs corresponding to different sensitivity ranges) might ease proper system configuration in such deployments.

2. Export Labeling:

Once the DOI is selected, the CALIPSO Sensitivity Label and values are determined based on the internal Sensitivity Label and the DOI. In the event the internal Sensitivity Level does not map to a valid CALIPSO Sensitivity Label, then an error SHOULD be returned to the upper-level protocol and that error MAY be logged. No further attempt to send this datagram should be made.

3. Access Control:

Once the datagram is marked and the sending logical interface is selected (by the routing code), the datagram's Sensitivity Label is compared against the Sensitivity Label range(s) associated with that logical interface. For the datagram to be sent, the interface MUST list the DOI of the datagram Sensitivity Label as one of the permissible DOI's and the datagram Sensitivity Label must be within range for the range associated with that DOI. If the datagram fails this access test, then

an error SHOULD be returned to the upper-level protocol and MAY be logged. No further attempt to send this datagram should be made.

6.2.2. Import

When a datagram arrives at an interface on an End System, the receiving End System MUST:

1. Verify the CALIPSO checksum. Datagrams with invalid checksums MUST be silently dropped. Such a drop event SHOULD be logged as a security fault with an indication of what happened.

2. Verify the CALIPSO has a known and valid DOI. Datagrams with unrecognized or illegal DOIs MUST be silently dropped. Such an event SHOULD be logged as a security fault with an indication of what happened.

3. Verify the DOI is a permitted one for the receiving interface. Datagrams with prohibited DOI values MUST be silently dropped. Such an event SHOULD be logged as a security fault with an indication of what happened.

4. Verify the CALIPSO Sensitivity Label is within the permitted range for the receiving interface:

   NOTE WELL: EACH permitted DOI on an interface has a separate table describing the permitted range for that DOI.

   A datagram with a Sensitivity Label within the permitted range is accepted for further processing.

   A datagram with a Sensitivity Label disjoint with the permitted range MUST be silently dropped. Such an event SHOULD be logged as a security fault, with an indication that the packet was dropped because of a disjoint Sensitivity Label. An ICMP error message MUST NOT be sent in this case.

   A datagram with a Sensitivity Label below the permitted range MUST be dropped. This event SHOULD be logged as a security fault, with an indication that the packet was below range. An ICMP error message MUST NOT be sent in this case.

A datagram with a Sensitivity Label above the permitted range MUST be dropped. This event SHOULD be logged as a security fault, with an indication that the packet was above range. An ICMP error message MUST NOT be sent in this case.

5. Once the datagram has been accepted, the receiving system MUST use the import Sensitivity Label and DOI to associate the appropriate internal Sensitivity Label with the data in the received datagram. This label information MUST be carried as part of the information returned to the upper-layer protocol.

6.3. Intermediate System Processing

This section describes CALIPSO-related processing for IPv6 packets transiting an IPv6 Intermediate System that claims to implement and comply with this specification. This document places no additional requirements on IPv6 Intermediate Systems that do not claim to comply or conform with this document.

The CALIPSO packet format has been designed so that one can configure an Intermediate System with the minimum sensitivity level, maximum Sensitivity Level, minimum compartment bitmap, and maximum compartment bitmap -- and then deploy that system without forcing the system to know the detailed human meaning of each Sensitivity Level or compartment bit value. Instead, once the minimum and maximum labels have been configured, the Intermediate System can apply a simple algorithm to determine whether or not a packet is within range for a given interface. This design should be straight-forward to implement in Application-Specific Integrated Circuit (ASIC) or Field Programmable Gate Array (FPGA) hardware, because the option format is simple and easy to parse, and because only a single comparison algorithm (defined in this RFC, hence known in advance) is needed.

6.3.1. Input

Intermediate Systems have slightly different rules for processing marked datagrams than do End Systems. Primarily, Intermediate Systems do not IMPORT or EXPORT transit datagrams, they just forward them. Also, in most deployments intermediate systems are used to provide Mandatory Access Controls to packets traversing more than one subnetwork.

The following checks MUST occur before any other processing. Upon receiving a CALIPSO-labeled packet, an Intermediate System must:

1. Determine whether or not this datagram is destined for (addressed to) this Intermediate System. If so, then the Intermediate System becomes an End System for the purposes of receiving this particular datagram and the rules for IMPORTing described above are followed.

2. Verify the CALIPSO checksum. Datagrams with invalid checksums MUST be silently dropped. The drop event SHOULD be logged as a security fault with an indication of what happened and MAY additionally be logged as a network fault.

   NOTE WELL:
   A checksum failure could indicate a general network problem (e.g., noise on a radio link) that is unrelated to the presence of a CALIPSO option, but it also could indicate an attempt by an adversary to tamper with the value of a CALIPSO label.

3. Verify the CALIPSO has a known and valid DOI. Datagrams with unrecognized or illegal DOIs MUST be silently dropped. Such an event SHOULD be logged as a security fault with an indication of what happened.

4. Verify the DOI is a permitted one for the receiving interface. Datagrams with prohibited DOIs MUST be silently dropped. Such a drop SHOULD be logged as a security fault with an indication of what happened.

5. Verify the Sensitivity Label within the CALIPSO is within the permitted range for the receiving interface:

   NOTE WELL:
   Each permitted DOI on an interface has a separate table describing the permitted range for that DOI.

   A rejected datagram with a Sensitivity Label below or disjoint with the permitted range MUST be silently dropped. Such an event SHOULD be logged as a security fault with an indication of what happened. An ICMP error message MUST NOT be sent in this case.

A rejected datagram with a Sensitivity Label above the permitted range MUST be dropped. The drop event SHOULD be logged as a security fault with an indication of what happened. An ICMP error message MUST NOT be sent in this case.

If and only if all the above conditions are met is the datagram accepted by the IPv6 Intermediate System for further processing and forwarding.

At this point, the datagram is within the permitted range for the Intermediate System, so appropriate ICMP error messages MAY be created by the IP module back to the originating End System regarding the forwarding of the datagram. These ICMP messages MUST be created with the exact same Sensitivity Label as the datagram causing the error. Standard rules about generating ICMP error messages (e.g., never generate an ICMP error message in response to a received ICMP error message) continue to apply. Note that these locally generated ICMP messages must go through the same outbound checks (including MAC checks) as any other forwarded datagram as described in the following paragraphs.

6.3.2. Translation by Intermediate Systems

It is at this point, after input processing and before output processing, that translation of the CALIPSO from one DOI to another DOI takes place in an Intermediate System, if at all. Section 6.4 describes the two possible approaches to translation.

6.3.3. Output

Once the forwarding code has selected the interface through which the datagram will be transmitted, the following takes place:

1. If the output interface requires that all packets contain a CALIPSO label, then verify that the packet contains a CALIPSO label.

2. Verify the DOI is a permitted one for the sending interface and that the datagram is within the permitted range for the DOI and for the interface.

3. Datagrams with prohibited DOIs or with out-of-range Sensitivity Labels MUST be dropped. Any drop event SHOULD be logged as a security fault, including appropriate details about which datagram was dropped and why.

4. Datagrams with prohibited DOIs or out-of-range Sensitivity Labels MAY result in an ICMP "Destination Unreachable" error message, depending upon the security configuration of the system.

   If the cause of the dropped packet is that the DOI is prohibited or unrecognized, then a reason code of "No Route to Host" is used. If the dropped packet's DOI is valid, but the Sensitivity Label is out of range, then a reason code of "Administratively Prohibited" is used. If an unlabeled packet has been dropped because the packet is required to be labeled, then a reason code of "Administratively Prohibited" is used.

   In all cases, if an ICMP Error Message is sent, then it MUST be sent with the same Sensitivity Label as the rejected datagram.

   The choice of whether or not to send an ICMP message, if sending an ICMP message for this case is implemented, MUST be configurable, and SHOULD default to not sending an ICMP message. Standard conditions about generating ICMP error messages (e.g., never send an ICMP error message about a received ICMP error message) continue to apply.

6.4. Translation

A system that provides on-the-fly relabeling is said to "translate" from one DOI to another. There are basically two ways a datagram can be relabeled:

Either the Sensitivity Label can be converted from a CALIPSO Sensitivity Label, to an internal Sensitivity Label, and then back to a new CALIPSO Sensitivity Label, exclusive-or a CALIPSO Sensitivity Label can be directly remapped into a new CALIPSO Sensitivity Label.

The first of these methods is the functional equivalent of "importing" the datagram then "exporting" it and is covered in detail in the "Import" (Section 6.2.2) and "Export" (Section 6.2.1) sections above.

The remainder of this section describes the second method, which is direct relabeling. The choice of which method to use for relabeling is an implementation decision outside the scope of this document.

A system that provides on-the-fly relabeling without importing or exporting is basically a special case of the Intermediate System rules listed above. Translation or relabeling takes place AFTER all input checks take place, but before any output checks are done.

Once a datagram has passed the Intermediate System input processing and input validation described in Section 6.3.1, and has been accepted as valid, the CALIPSO in that datagram may be relabeled. To determine the new Sensitivity Label, first determine the new output DOI.

The selection of the output DOI may be based on any of Incoming DOI, Incoming Sensitivity Label, Destination End System, Destination Network, Destination Subnetwork, Sending Interface, or Receiving Interface, or combinations thereof. Exact details on how the output DOI is selected are implementation dependent, with the caveat that it should be consistent and reversible. If a datagram from End System A to End System B with DOI X maps into DOI Y, then a datagram from B to A with DOI Y should map into DOI X.

Once the output DOI is selected, the output Sensitivity Label is determined based on (1) the input DOI and input Sensitivity Label and (2) the output DOI. In the event the input Sensitivity Label does not map to a valid output Sensitivity Label for the output DOI, then the datagram MUST be silently dropped and the drop event SHOULD be logged as a security fault.

Once the datagram has been relabeled, the Intermediate System output procedures described in Section 6.3.3 are followed, with the exception that any error that would cause an ICMP error message to be generated back to the originating End System instead MUST silently drop the datagram without sending an ICMP error message. Such a drop SHOULD be logged as a security fault.

7. Architectural and Implementation Considerations

This section contains "implementation considerations"; it does not contain "requirements". Implementation experience might eventually turn some of them into implementation requirements in some future version of this specification.

This IPv6 option specification is only a small part of an overall distributed Multi-Level Secure (MLS) deployment. Detailed instructions on how to build a Multi-Level Secure (MLS) device are well beyond the scope of this specification. Additional information on implementing a Multi-Level Secure operating system, for example implementing an MLS End System, is available from a range of sources [TCSEC] [TNI] [CMW] [CC] [ISO-15408] [MLOSPP].

Because the usual 5-tuple (i.e., Source IP address, Destination IP address, Transport protocol, Source Port, and Destination Port) do not necessarily uniquely identify a flow within a labeled MLS network deployment, some applications or services might be impacted by multiple flows mapping to a single 5-tuple. This might have unexpected impacts in a labeled MLS network deployment using such application protocols. For example, Resource Reservation Protocol (RSVP), Session Initiation Protocol (SIP), and Session Description Protocol (SDP) might be impacted by this.

A number of Commercial-Off-The-Shelf (COTS) applications (e.g., Remote Access Dial-In User Service (RADIUS), Hyper-Text Transfer Protocol (HTTP), and Transport-Layer Security (TLS) web content access) have been included in MLS network deployments for about two decades, without operational difficulties or a need for special modifications. The ability to use these common applications demonstrates that the basic Internet architecture remains unchanged in an MLS deployment, although certain details (e.g., adding labels to IP datagrams) do change.

7.1. Intermediate Systems

Historically, RFC 1108 was supported by one commercial label-aware IP router. Neither RFC 1038 nor FIPS-188 were supported in any commercial IP router, so far as the authors are aware. A label-aware router does not necessarily use an MLS operating system. Instead, a label-aware router might use a conventional router operating system, adding extensions to permit application of per-logical-interface label-oriented Access Control Lists (ACLs) to IP packets entering and leaving that router’s network interface(s).

This proposal does not change IP routing in any way. Existing label-aware routers do not use Sensitivity Labels in path calculations, Routing Information Base (RIB) or Forwarding Information Base (FIB) calculations, their routing protocols, or their packet forwarding decisions.

Similarly, existing MLS network deployments use many protocols or specifications, for example, Differentiated Services, without modification. For Differentiated Services, this might mean that multiple IP flows (i.e., flows differing only in their CALIPSO label value) would be categorized and handled by Intermediate Systems as if they were a single flow.

Router performance is optimized if there is hardware support for applying the Mandatory Access Controls based on this label option. An issue with CIPSO is that the option syntax is remarkably complex [FIPS-188]. So this label option uses a simplified syntax. This

should make it more practical to create custom logic (e.g., in Verilog) with support for this option and the associated Mandatory Access Controls.

7.2. End Systems

It is possible for a system administrator to create two DOIs with different overlapping compartment ranges. This can be used to reduce the size of the IPv6 Sensitivity Label option in some deployments.

7.3. Upper-Layer Protocols

As CALIPSO is an IP option, this document focuses upon the network-layer handling of IP packets containing CALIPSO options. This section provides some discussion of some upper-layer protocol issues.

This section is not a complete specification for how an MLS End System handles information internally after the decision has been made to accept a received IPv6 packet containing a CALIPSO option. Implementers of MLS systems might wish also to consult [TCSEC], [TNI], [CMW], [CC], [ISO-15408], and [MLOSPP].

In a typical MLS End System, the information received from the network (i.e., information not dropped by the network layer as a result of the CALIPSO processing described in this document) is assigned an internal Sensitivity Label while inside the End System operating system. The MLS End System uses the Bell-LaPadula Mandatory Access Control policy [BL73] to determine how that information is processed, including to which transport-layer sessions or to which applications the information is delivered.

Within this section, we use one additional notation, in an attempt to be both clear and concise. Here, the string "W:XY" defines a Sensitivity Label where the Sensitivity Level is W and where X and Y are the only compartments enabled, while the string "W::" defines a Sensitivity Label where the Sensitivity Level is W and there are no compartments enabled.

7.3.1. TCP-Related Issues

With respect to a network, each distinct Sensitivity Label represents a separate virtual network, which shares the same physical network.

The above statement, taken from Section 3, has a non-obvious, but critical, corollary. If there are separate virtual networks, then it is possible for a system that exists in multiple virtual networks to have identical TCP connections, each one existing in a different virtual network.

TCP connections are normally identified by source and destination port, and source and destination address. If a system labels datagrams with the CALIPSO option (which it must do if it exists in multiple virtual networks - e.g., a "Multi-Level Secure" system), then TCP connections are identified by source and destination port, source and destination address, and an internal Sensitivity Label (optionally, a Sensitivity Label range). This corrects a technical error in RFC 793, and is consistent with all known MLS operating system implementations [TNI] [RFC793]. There are no known currently deployed TCP instances that actually comply with this specific detail of RFC 793.

7.3.2. UDP-Related Issues

Unlike TCP or SCTP, UDP is a stateless protocol, at least conceptually. However, many implementations of UDP have some session state (e.g., Protocol Control Blocks in 4.4 BSD), although the UDP protocol specifications do not require any state.

One consequence of this is that in widely used End System implementations of UDP and IPv6, a UDP listener might be bound only to a particular UDP port on its End System -- without binding to a particular remote IP address or local IP address.

UDP can be used with unicast or with multicast. Some existing UDP End System implementations permit a single UDP packet to be delivered to more than one listener at the same time. Except for the application of Mandatory Access Controls, the behavior of a given system should remain the same (so that application behavior does not change in some unexpected way) with respect to delivery of UDP datagrams to listeners.

For example, if a listener on UDP port X has a Sensitivity Label range with a minimum of "S:AB" and a maximum of "S:AB", then only datagrams with a destination of UDP port X and a Sensitivity Label of "S:AB" will be delivered to that listener.

For example, if a listener on UDP port Y has a Sensitivity Label range with a minimum of "W::" and a maximum of "X:ABC" (where X dominates W), then a datagram addressed to UDP port Y with a Sensitivity Label of "W:A" normally would be delivered to that listener.

7.3.3. SCTP-Related Issues

With respect to a network, each distinct Sensitivity Label represents a separate virtual network, which shares the same physical network.

The above statement, taken from Section 3, has a non-obvious, but critical, corollary. If there are separate virtual networks, then it is possible for a system that exists in multiple virtual networks to have identical SCTP connections, each one existing in a different virtual network.

As with TCP, SCTP is a connection-oriented transport protocol and has substantial session state. Unlike TCP, SCTP can support session-endpoint migration among IP addresses at the same end node(s), and SCTP can also support both one-to-one and one-to-many communication sessions.

In single-level End Systems, in the one-to-one mode, the SCTP session state for a single local SCTP session includes the set of remote IP addresses for the single remote SCTP instance, the set of local IP addresses, the remote SCTP port number, and the local SCTP port number.

In single-level End Systems, in the one-to-many mode, the SCTP session state for a single local SCTP instance can have multiple concurrent connections to several different remote SCTP peers. There cannot be more than one connection from a single SCTP instance to any given remote SCTP instance. Thus, in single-level End Systems, in the one-to-many mode, the local SCTP session state includes the set of remote IP addresses, the set of local IP addresses, the remote SCTP port number for each remote SCTP instance, and the (single) local SCTP port number.

In MLS End Systems, for either SCTP mode, the SCTP session state additionally includes the Sensitivity Label for each SCTP session. A single SCTP session, whether in the one-to-one mode or in the one-to-many mode, MUST have a single Sensitivity Label, rather than a Sensitivity Label range.

Unlike TCP, SCTP has the ability to shift an existing SCTP session from one endpoint IP address to a different IP address that belongs to the same endpoint, when one or more endpoints have multiple IP addresses. If such shifting occurs within an MLS deployment, it is important that it only move to an IP address with a Sensitivity Label range that includes that SCTP session’s own Sensitivity Label.

Further, although a node might be multi-homed, it is entirely possible that only one of those interfaces is reachable for a given Sensitivity Label value. For example, one network interface on a node might have a Sensitivity Label range from "A::" to "B:XY" (where B dominates A), while a different network interface on the same node might have a Sensitivity Label range from "U::" "U::" (where A dominates U). In that example, if a packet has a CALIPSO label of

"A:X", then that packet will not be able to use the "U"-only network interface. Hence, an SCTP implementation needs to consider the Sensitivity Label of each SCTP instance on the local system when deciding which of its own IP addresses to communicate to the remote SCTP instance(s) for that SCTP instance. This issue might lead to novel operational issues with SCTP sessions. Implementers ought to give special attention to this SCTP-specific issue.

7.3.4. Security Logging

This option is recommended for deployment only in well-protected private networks that are NOT connected to the global Internet. By definition, such private networks are also composed only of trusted systems that are believed to be trustworthy. So the risk of a denial-of-service attack upon the logging implementation is much lower in the intended deployment environment than it would have been for general Internet deployments.

8. Security Considerations

This document describes a mechanism for adding explicit Sensitivity Labels to IPv6 datagrams. The primary purpose of these labels is to facilitate application of Mandatory Access Controls (MAC) in End Systems or Intermediate Systems that implement this specification.

As such, correct implementation of this mechanism is very critical to the overall security of the systems and networks where this mechanisms is deployed. Use of high-assurance development techniques is encouraged. End users should carefully consider the assurance requirements of their particular deployment, in the context of that deployment’s prospective threats.

A concern is that since this label is used for Mandatory Access Controls, some method of binding the Sensitivity Label option to the rest of the packet is needed. Without such binding, malicious modification of the Sensitivity Label in a packet would go undetected. So, implementations of this Sensitivity Label option MUST also implement support for the IP Authentication Header (AH). Implementations MUST permit the system administrator to configure whether or not AH is used.

ESP with null encryption mechanism can only protect the payload of an IPv6 packet, not any Hop-by-Hop Options. By contrast, AH protects all invariant headers and data of an IPv6 packet, including the CALIPSO Hop-by-Hop Option. The CALIPSO option defined in this document is always an IPv6 Hop-by-Hop Option, because the CALIPSO option needs to be visible to, and parsable by, IPv6 routers and security gateways so that they can apply MAC policy to packets.

It is anticipated that if AH is being used with a symmetric authentication algorithm, then not only the recipient End System, but also one or more security gateways along the path, will have knowledge of the symmetric key -- so that AH can be used to authenticate the packet, including the CALIPSO label. In this case, all devices knowing that symmetric authentication key would need to be trusted. Alternatively, AH may be used with an asymmetric authentication algorithm, so that the recipient and any security gateways with knowledge of the authentication key can authenticate the packet, including the CALIPSO label.

If AH or ESP are employed to provide "labeled IP Security" within some CALIPSO deployment, then the Sensitivity Label of the IP Security Association used for a given packet MUST have the same meaning as the Sensitivity Label carried in the CALIPSO option of that packet, in order that MAC policy can and will be correctly applied.

Because the IP Authentication Header will include the CALIPSO option among the protected IPv6 header fields, modification of a CALIPSO-labeled packet that also contains an IP Authentication Header will cause the resulting packet to fail authentication at the destination node for the AH security session. Therefore, CALIPSO labels cannot be inserted, deleted, or translated for IPv6 packets that contain an IP Authentication Header.

NOTE WELL: The "not modified during transit" bit for IPv6 option types was really created to be the "include in AH calculations" signal. There was no other reason to define that bit in IPv6.

In situations where a modification by an Intermediate System is required by policy, but is not possible due to AH, then the packet MUST be dropped instead. If the packet must be dropped for this reason, then an ICMP "Destination Unreachable" error message SHOULD be sent back to the originator of the dropped packet with a reason code of "Administratively Prohibited". If the packet can be forwarded properly without violating the MLS MAC policy of the Intermediate System, then (by definition) such a packet modification is not required.

Note that in a number of error situations with labeled networking, an ICMP error message MUST NOT be sent in order to avoid creating security problems. In certain other error situations, an ICMP error message might be sent. Such ICMP handling details have been described earlier in this document. Even if an ICMP error message is sent, it might be dropped along the way before reaching its intended destination -- due to MAC rules, DOI differences, or other configured security policies along the way from the node creating the ICMP error

message to the intended destination node. In turn, this can mean operational faults (e.g., fibre cut, misconfiguration) in a labeled network deployment might be more difficult to identify and resolve.

This mechanism is only intended for deployment in very limited circumstances where a set of systems and networks are in a well-protected operating environment and the threat of external or internal attack on this mechanism is considered acceptable to the accreditor of those systems and networks. IP packets containing visible packet labels ought never traverse the public Internet.

This specification does not seek to eliminate all possible covert channels. The TCP specification clarification in Section 7.3.1 happens to reduce the bandwidth of a particular known covert channel, but is present primarily to clarify how networked MLS systems have always been implemented [TNI] [MLOSPP].

Of course, subject to local security policies, encrypted IPv6 packets with CALIPSO labels might well traverse the public Internet after receiving suitable cryptographic protection. For example, a CALIPSO-labeled packet might travel either through a Tunnel-mode ESP (with encryption) VPN tunnel that connects two or more MLS-labeled network segments. Alternatively, a CALIPSO-labeled IPv6 packet might travel over some external link that has been protected by the deployment of evaluated, certified, and accredited bulk encryptors that would encrypt the labeled packet before transmission onto the link and decrypt the labeled packet after reception from the link.

Accreditors of a given CALIPSO deployment should consider not only personnel clearances and physical security issues, but also electronic security (e.g., TEMPEST), network security (NETSEC), communications security (COMSEC), and other issues. This specification is only a small component of an overall MLS network deployment.

9. IANA Considerations

9.1. IP Option Number

An IPv6 Option Number [RFC2460] has been registered for CALIPSO.

HEX	BINARY			
	act	chg	rest	
---	---	---	---	---
7	00	0	00111	CALIPSO

For the IPv6 Option Number, the first two bits indicate that the IPv6 node skip over this option and continue processing the header if it does not recognize the option type. The third bit indicates that the Option Data must not change en route.

This document is listed as the reference document.

9.2. CALIPSO DOI Values Registry

IANA has created a registry for CALIPSO DOI values. The initial values for the CALIPSO DOI registry, shown in colon-separated quad format, are as follows:

DOI Value	Organization or Use
0:0:0:0	NULL DOI. This ought not be used on any network.
0:0:0:1 to 0:255:255:255	For private use among consenting parties within private networks.
1:0:0:0 to 254:255:255:255	For assignment by IANA to organizations following the Expert Review procedure [RFC5226].
255:0:0:0 to 255:255:255:255	Reserved to the IETF for future use by possible revisions of this specification.

The CALIPSO DOI value 0:0:0:0 is the NULL DOI and is not to be used on any network or in any deployment.

All other CALIPSO DOI values beginning with decimal 0: are reserved for private use amongst consenting parties; values in this range will not be allocated by IANA to any particular user or user community.

For the CALIPSO DOI values 1:0:0:0 through 254:255:255:255 (inclusive), IANA should follow the Expert Review procedure when DOI Allocation requests are received.

CALIPSO DOI values beginning with decimal 255 are reserved to the IETF for potential future use in revisions of this specification. IESG approval is required for allocation of DOI values within that range.

10. Acknowledgments

This document is directly derived from an Internet-Draft titled "Son of IPSO (SIPSO)" written by Mike StJohns circa 1992. Various changes have been made since then, primarily to support IPv6 instead of IPv4. The concepts, most definitions, and nearly all of the processing rules here are identical to those in that earlier document.

Steve Brenneman, L.C. Bruzenak, James Carlson, Pasi Eronen, Michael Fidler, Bob Hinden, Alfred Hoenes, Russ Housley, Suresh Krishnan, Jarrett Lu, Dan McDonald, Paul Moore, Joe Nall, Dave Parker, Tim Polk, Ken Powell, Randall Stewart, Bill Sommerfeld, and Joe Touch (listed in alphabetical order by family name) provided specific feedback on earlier versions of this document.

The authors also would like to thank the several anonymous reviewers for their feedback, and particularly for sharing their insights into operational considerations with MLS networking.

The authors would like to thank the IESG as a whole for providing feedback on earlier versions of this document.

11. References

11.1. Normative References

[RFC1662] Simpson, W., Ed., "PPP in HDLC-like Framing", STD 51, RFC 1662, July 1994.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

11.2. Informative References

[BL73] Bell, D.E. and LaPadula, L.J., "Secure Computer Systems: Mathematical Foundations and Model", Technical Report M74-244, MITRE Corporation, Bedford, MA, May 1973.

[CW87] D.D. Clark and D.R. Wilson, "A Comparison of Commercial and Military Computer Security Policies", in Proceedings of the IEEE Symposium on Security and Privacy, pp. 184-194, IEEE Computer Society, Oakland, CA, May 1987.

[CMW] US Defense Intelligence Agency, "Compartmented Mode Workstation Evaluation Criteria", Technical Report DDS-2600-6243-91, Washington, DC, November 1991.

[DoD5200.1-R] US Department of Defense, "Information Security Program Regulation", DoD 5200.1-R, 17 January 1997.

[DoD5200.28] US Department of Defense, "Security Requirements for Automated Information Systems," Directive 5200.28, 21 March 1988.

[MLOSPP] US Department of Defense, "Protection Profile for Multi-level Operating Systems in Environments requiring Medium Robustness", Version 1.22, 23 May 2001.

[ISO-15408] International Standards Organisation, "Evaluation Criteria for IT Security", ISO/IEC 15408, 2005.

[CC] "Common Criteria for Information Technology Security Evaluation", Version 3.1, Revision 1, CCMB-2006-09-001, September 2006.

[TCSEC] US Department of Defense, "Trusted Computer System Evaluation Criteria", DoD 5200.28-STD, 26 December 1985.

[TNI] (US) National Computer Security Center, "Trusted Network Interpretation (TNI) of the Trusted Computer System Evaluation Criteria", NCSC-TG-005, Version 1, 31 July 1987.

[FIPS-188] US National Institute of Standards and Technology, "Standard Security Labels for Information Transfer", Federal Information Processing Standard (FIPS) 188, September 1994.

[IEEE802.1Q] IEEE, "Virtual Bridged Local Area Networks", IEEE Standard for Local and metropolitan area networks, 802.1Q - 2005, ISBN 0-7381-4876-6, IEEE, New York, NY, USA, 19 May 2006.

[RFC791] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.

[RFC793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, September 1981.

[RFC1038] St. Johns, M., "Draft revised IP security option", RFC 1038, January 1988.

[RFC1108] Kent, S., "U.S. Department of Defense Security Options for the Internet Protocol", RFC 1108, November 1991.

[RFC1825] Atkinson, R., "Security Architecture for the Internet Protocol", RFC 1825, August 1995.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.

[RFC4302] Kent, S., "IP Authentication Header", RFC 4302, December 2005.

[RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, December 2005.

Authors' Addresses

Michael StJohns
Germantown, MD
USA

EMail: mstjohns@comcast.net

Randall Atkinson
Extreme Networks
3585 Monroe Street
Santa Clara, CA
USA 95051

EMail: rja@extremenetworks.com
Phone: +1 (408)579-2800

Georg Thomas
US Department of Defense
Washington, DC
USA