

# Security Label-based Authorization in Directory Services

Steve Kille  
Kurt D. Zeilenga

Isode Limited

## Abstract

Security Labels provide an important mechanism for controlling access to information in many high security environments, and are also useful in environments with lower security requirements. This paper provides a reasonably detailed description of how security labels and clearances work, both in general and in LDAP/X.500 directory services, while attempting to avoid some of the significant complexity often attributed to the subject.

## Introduction

Security Labels provide an important mechanism for controlling access to information in many high security environments, and are also useful in environments with lower security requirements. This paper provides a reasonably detailed description of how security labels and clearances work, while attempting to avoid the high level of technical complexity seen in many papers in this area.

This paper starts by looking at how security labeling is used for paper documents and other non-electronic applications. Then it looks at how Security Labels work for electronic documents and other online applications.


## How (non-electronic) Security Labels Work

The basic mechanism of security labels is familiar to most people. Documents are labeled with a classification, such as "Confidential", "Secret", or "Top Secret". This security label will be clearly visible on the document.

People are given a clearance, using the same scheme. For example, someone may be cleared to "Secret" level, meaning that they can read documents labeled "Secret", but not a document labeled "Top Secret".

Security labels on documents are just one aspect of the model, as a security label can apply to any information. In discussions, generic information such as the role of a person can be labeled as "Secret". A key concept is that security labels are associated with information.

The core model of security clearance is that a person (i.e., someone accessing information) has a security clearance, that controls what information can be accessed. Thus, access to information is controlled by matching the security clearance of the person accessing information with the security label associated with the information. This is illustrated above. Bill has a document that is labelled "Top Secret". He will check clearances of "Joe" and "Fred". He will give the document to Joe, who has "Top Secret" clearance, but not to "Fred", who only has "Secret" clearance.



Security clearance may also be associated with a location. For example, if a meeting room is cleared to "Secret" level, it may be used for discussions of "Secret" information, but may not be used for discussions of "Top Secret" information. This follows the basic model that security labels are matched against security clearance.

## Why Security Labels are Used

Security labels are widely used as a mechanism for controlling information access, for a number of reasons.

1. The model of security labels and clearances is very easy to understand. This is important, as complex models are more prone to user error.
2. Clearances can be managed for large numbers of people, employed by different organizations and in different locations.
3. Other schemes would not be practical, particularly where there are large numbers of people involved. Schemes that would not be practical in general:
  - a) Listing on each document, who can access it.
  - b) Record for each user, which documents they can access.
  - c) Maintain independent lists, to correlate user and document access.

The summary is that security labeling is a practical scheme, that has been used of many years.

# What is in a Security Label

The most important part, and often the only part, of a security label is the classification. Many government security label schemes use the following classifications:

- Unclassified
- Restricted
- Confidential
- Secret
- Top Secret

This is an ordered scheme, so that person cleared to Secret level, can access Secret, Confidential, Restricted, and Unclassified information.

Other approaches to classification are possible. For example, as described in RFC 3114, Amoco use the classifications "Amoco-general", "Amoco-confidential" and "Amoco-highly-confidential".

In many situations, the classification does not give a sufficiently fine grain of control. To deal with this, security labels may contain, in addition to the classification, additional information know as categories. Some examples:

1. Security labels are often used to control information related to national security. A category of country is often used, with users assigned category value of their nationality. Information is then controlled by country (e.g., "two eyes" for release to two specific countries) using the category to control where the information is released to.
2. An agency controls information by topic. Here a category "Biological Weapons" could be used to restrict access to those cleared for this category. In order to access this information, a user needs to be cleared to the right classification and be cleared for the category.

There are three specific types of category:

1. Restrictive. Here the user must have clearance for all values of the category set in the label. This is useful to apply a number of additional controls.
2. Permissive. Here the user must have a clearance for one of the categories set in the label. This could be used where information is cleared for several countries (indicated in the label) and a user needs to be cleared for at least one of these.
3. Informative. The category information in the label is made available to the user, but is not checked against clearance.

The categorization itself may be classified, which means that the category values may only be shared between those cleared at the appropriate classification.

## Checking Security Labels against Clearance

A document or information with a given security label will start off being held or created by someone with appropriate clearance. The key checks will happen when that document or information is transferred to another person. If the other person has an appropriate clearance, the information may be transferred.

Matching security label and security clearance is straightforward. The other part of the problem is for the person providing information to determine the security clearance of the recipient. In many high security organizations, employees will wear photo-id badges with the person's clearance clearly written on it (or implicit from a color code). A visitor's clearance will have been set by an organization trusted by the organization being visited and held by that organization. Prior to the visitor arriving, the security office of the organization being visited will verify the visitor's clearance with the organization that holds this clearance. On the visit, the visitor's identity will be verified, and a badge indicating the security clearance issued. The details of this process will vary, but it can be seen that verifying security clearance is the most complex part of the process.

## **Security Policy**

Security policy is a generic term, but in the context of access control using security labels and security clearances it primarily relates to two things:

1. The Security Policy defines the security label values that are valid.
2. The Security Policy defines how security labels are matched against security clearance.

The term "Security Policy" will be used in this document to refer to this quite narrow definition, and "security policy" where the term is used more generically. Security Policy will be part of a much broader security policy, which will cover such things as rules for assigning security labels and process for vetting people for security clearance.

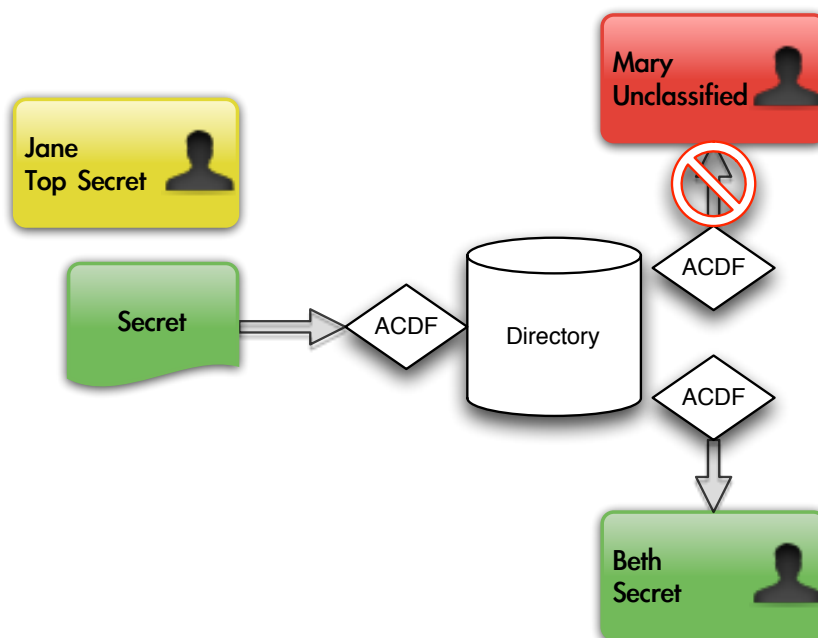
Security Policy will vary between organizations. Different governments and organizations will have different Security Policies.

## **Why Electronic Security Labels are Useful**

A major requirement for electronic security labels is to support organizations that use security labels and security clearances. Security labels apply to all information. If information is going to be communicated electronically the same security label and security clearance based access control should be used. When communication is electronic, it is desirable for the communication system to enforce Security Policy, and ensure that information is not sent to someone that does not have an appropriate clearance. The primary consideration on electronic security labels is for support of this sort of organization.

Security Labels and Security Clearance provides a mechanism for controlling access to information that works well for large numbers of users. It can be an effective approach for access control in organizations that do not use non-electronic security labels.

The following diagram shows the translation of the earlier example into electronic form. Here, Jane, who has a "Top Secret" clearance, adding a directory object labeled "Secret". As her clearance includes "Secret", this operation is allowed. Beth has "Secret" clearance and so may read the object. Mary has "Unclassified" clearance and so may not read the object.



Access control checking based on Security Labels needs to occur in several places. While the above

diagram shows checks done by a directory service, other entities (such as each user's directory client) may perform checks. However, it is common for the directory applications for the directory service to be access control enforcement entity.

## Electronic Representation of Security Label

### Requirements of an Online Representation

The central part of an electronic security label scheme is the format used to represent security labels. An online representation needs to deal with a number of things:

1. It should be useable with a wide range of protocols and formats, including documents, email, instant messaging, directory data, and database data.
2. It should support a wide range of Security Policies, so that it does not restrict the organizations that can use it.
3. It should be compact. Reasons:
  - Some applications will need to label quite small pieces of information, and addition of Security Labels should not be too high.
  - Some organizations using Security Labels, particularly military, operate in low bandwidth situation.
4. It should integrate well with digital signatures.

### Electronic Representation of Security Labels

Security Labels are generally specified by the standard that uses them. There are two compatible specifications of Security Labels in wide use:

1. X.400 Messaging and X.500 Directory use a common definition, specified in X.411 and X.501 respectively.
2. SMTP Message uses ESS Security Labels as defined in the "Enhanced Security Services for S/MIME" (RFC 2634), which differ only slightly from X.400/X.500 security labels.

The rest of this document refers to Security Labels, as online representations conformant to these specifications.

Object identifiers are a compact representation of unique values, based on an internationally allocated hierarchy of numbers. It is straightforward for any organization to get a part of the hierarchy and to allocate further values. For example, the US DoD has the object identifier value: 1.3.6. Object identifiers are used in many protocols and are important for security labels.

A Security Label has the following components:

- Security Policy. This is an object identifier that identifies the policy, and will be a value allocated by the organization setting the policy. This gives a compact mechanism to represent one or more policies set by an organization or government.
- Classification. The security classification is represented by an integer. Six integer values represent standard classifications (unmarked; unclassified; restricted; confidential; secret; top secret). The semantics of other values of the classification are defined by Security Policy.
- Categories. A Security Label may have one or more category value. The syntax and semantics of a category may vary with Security Policy. In order to achieve this a category value consists of an object identifier, and data whose syntax and semantics are defined by that object identifier.
- Privacy Mark. A string value used in certain authorization systems (generally not used).

This structure includes all of the necessary information in a compact encoding. A key feature of this representation is that it is extensible, and can be used to represent a wide variety of Security Policies.

A Security Label will be associated with the object that is being labeled. For a document or message, the Security Label will be included within the document or message.

## **Electronic Representation of Security Clearance**

Security Clearance is defined in X.501, and has a structure that corresponds to Security Label. A Security Clearance comprises the following elements.

- Security Policy. An object identifier, as in security label.
- Categories. As in security label.
- Classification List. This is a list of classifications, represented as a bit string, with bit values corresponding to the integer values of Classification in security label. The reason for this, is that Classifications are not always ordered (although they often are). Because of this, the clearance needs to represent all of the Classification values for which the holder of the Security Clearance is cleared.

It can be seen that with this very similar structure, that it is straightforward to perform basic matching of Security Label and Security Clearance.

A Security Clearance is associated with a user typically through a X.509 certificate or attribute certificate. Another, more simple (but less secure) approach to handling this is to put the user's Security Clearance in an attribute of the user's directory entry. This makes it straightforward to determine the clearance of a user.

## **Electronic Representation of Security Policy**

Security Policy is a very important part of handling Security Labels and Security Clearances, and is key to use of a single implementation by different organizations with different Security Policies. A user with a clearance from the French Government will not (usually) get access to documents classified by the UK Government. The policy of Security Label and Security Clearance need to match, and this component is fundamental to ensuring separation.

Security Policy is also important for defining which classification and category values are used, and can cover wide range of functionality. A detailed examination of Security Policy will be the subject of a future Isode white paper.

In order to support multiple security policies, a good implementation will have an external electronic representation of Security Policy so that this information can be shared and updated across all participating systems.

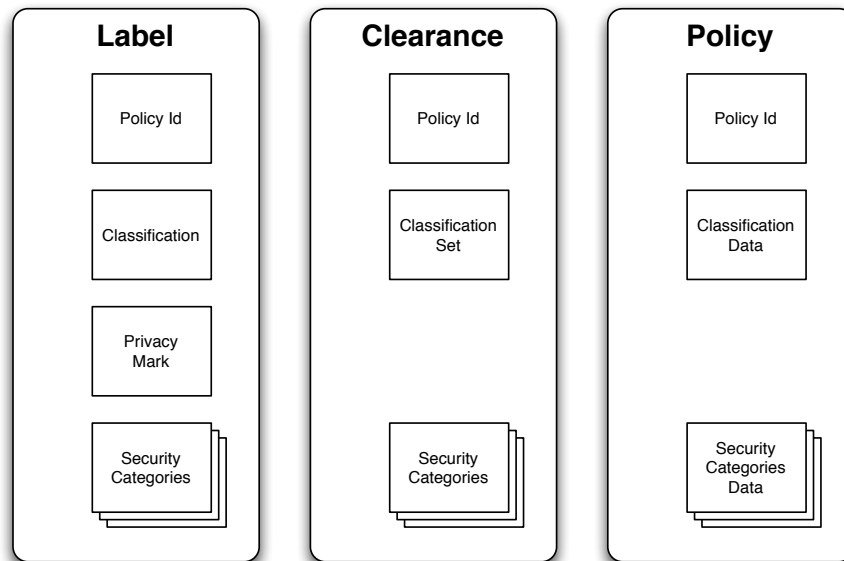
There are two standard definitions of representing Security Policy in Security Policy Information File (SPIF):

- X.841. "Security techniques - Security information objects for access control", published by the ITU (International Telecommunications Union).
- SDN.801. "Access control concept and mechanisms", published by the US National Security Agency.

These two standards have broadly similar capabilities, but are not strictly compatible. The benefit of using a standardized SPIF is that it enables Security Policy information to be shared between implementations from different vendors.


# Security Information Objects

The following diagram compares the Security Label, Clearance, and Policy structures:



## Access Control Decision Function

The following diagram provides an overview of the Access Control Decision Function used in Security Label-based Authorization:



## Integration with Digital Signatures

Security labels and security clearances will often be used by organizations with stringent overall security requirements, and it will often be essential to use digital signatures in conjunction with security labels and security clearances.

The most important use of digital signatures is to bind a Security Clearance to a user, in order to verify that the user has the Security Clearance claimed (as opposed to reading it from a directory entry and trusting the directory). X.509 digital certificates can include an X.501 Security Clearance, and so this integration is clean and straightforward. X.509 digital certificates can also include multiple name forms for a user (e.g., and Internet email address), and so the

certificate can be used to bind (using a digital signature) a Security Clearance to a name other than a directory name.

A Security Label will be associated with a piece of information, such as a directory object. If a digital signature is used, it will need to apply to both the Security Label and the labeled object; Their needs to be equivalent protection to tampering with either the label or the object.

## **Use in Directories**

X.500 specifications provide an attribute-value level security label-based access control system based upon attribute value contexts.

As LDAP does not include support for attribute value contexts, the X.500 security label-based access control system cannot be in LDAP. Isode has designed a simple access control system supporting entry and attribute level labels for use in X.500 and LDAP, and currently offers the former in its M-Vault directory server product.

## **Conclusions**

Security labels and security clearances are an important access control mechanism for many organizations. This can be supported electronically for a range of applications using Security Labels, and it can be integrated with digital signatures.