Access Control using Security Labels & Security Clearance

Isode

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 labelling 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 Security Labels Work (non-electronic)

The basic mechanism of security labels is familiar to most people. Documents are labelled 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 "Secret", but not a document labelled "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 labelled as "Secret".

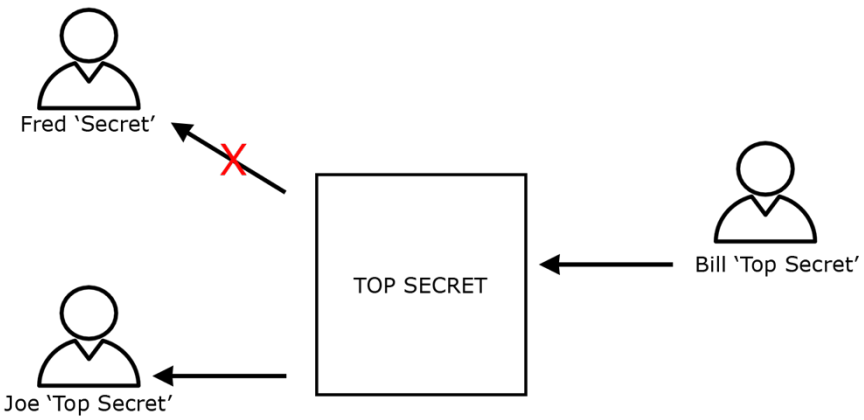

The core model of security clearance is that a person (i.e., someone accessing information) has a security clearance that controls what information can be accessed. Thus, access to information is controlled by matching the security clearance of the person accessing information with the security label 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".

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.

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:
    - Listing on each document, who can access it.
    - Record for each user, which documents they can access.
    - Maintain independent lists, to correlate user and document access.

The summary is that security labelling 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 someone cleared to Secret level, can access Secret, Confidential, Restricted, and Unclassified information.

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

Classifications

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 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, Bill has an online document labelled "Top Secret". Joe has "Top Secret" clearance (and so may receive the document) "Secret" clearance (and so may not receive the document).

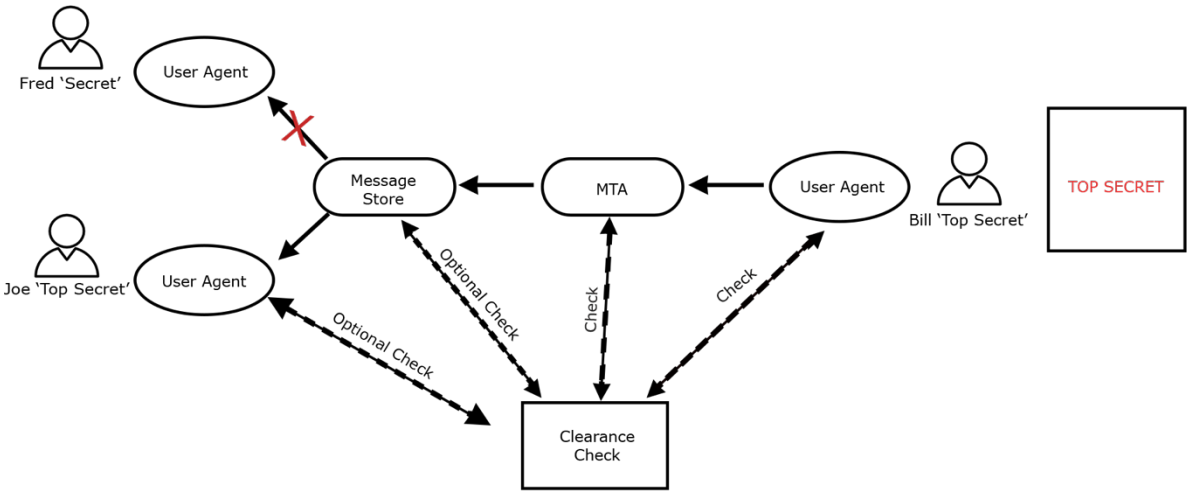

Access control checking based on Security Labels needs to occur in several places:

1. In Bill's User Agent (UA). Here, the UA should look up the intended recipients, and check the document Security Label against the recipient's Security Clearance. This will ensure that the UA only sends documents appropriately.
2. The Message Transfer Agent should also perform checks, as in general the communications infrastructure should not rely on correct client behavior. It should only deliver information where the Security Clearance matches the Security Label.
3. The Message Store may also perform checks before accepting a message for a user.

4. The receiving UA may also perform checks, so that a user will only be shown appropriate messages. This may be desirable, but should not be relied on.

It is important to note that access control decisions based on Security Labels will happen in a number of places.

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 usable 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 three major specifications of Security Labels:

1. NATO has standardized labels in STANAG 4774. Details in the Isode white paper Isode Approach to Data Centric Security using NATO Confidentiality Labels. This is expected to become the dominant label format for military use.
2. ESS Security Labels are defined in the Internet Standard "Enhanced Security Services for S/MIME" (RFC 2634). ESS Security Labels are used in S/MIME (the Internet Standard for secure documents and email) and within STANAG 4406 (military messaging).
3. X.411 and X.500 directory use a definition that is very similar to ESS, and is used by X.500 directory and by X.400 messaging.

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.

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 labelled. For a document or message, the Security Label will be included within the document or message.

Electronic Representation of Security Clearance

Security Clearances are standardized in STANAG 4774 with an older standard in X.501, and have 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. A simple 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 older standard definitions of representing Security Policy in Security Policy Information File (SPIF). These standards are:

- 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 compatible.

Isode and other vendors have moved to an open SPIF standard: Open XML SPIF, which is broadly based on SDN.801. It is anticipated that this will be the basis for a NATO standard SPIF.

The benefit of using a standardized SPIF is that it enables Security Policy information to be shared between implementations from different vendors.

More details on SPIFs are provided in the Isode white paper Creating and Managing a Security Label Policy.

Access Control

Access control using security labels is formally checking a security label against a security clearance in context of a SPIF. This is key for a wide range of applications and described for Isode applications in two white papers:

- Security Label Capabilities in M-Switch Products
- Using Security Labels to Control Message Flow in XMPP Services

NATO STANAG 4474 calls this access control Confidentiality Metadata Based Access Control (CMBAC – pronounced “come back”).

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.

Isode

www.isode.com

14 Castle Mews, Hampton
Middlesex, TW12 2NP

Secure, Seamless Communication Solutions