                 Implementing Company Classification Policy
                      with the S/MIME Security Label

Status of this Memo

Copyright Notice

Abstract

   This document discusses how company security policy for data
   classification can be mapped to the S/MIME security label. Actual
   policies from three companies provide worked examples.

1. Introduction

   Security labels are an optional security service for S/MIME. A
   security label is a set of security information regarding the
   sensitivity of the content that is protected by S/MIME encapsulation.
   A security label can be included in the signed attributes of any
   SignedData object. A security label attribute may be included in
   either the inner signature, outer signature, or both. The syntax and
   processing rules for security labels are described in RFC 2634 [ESS].

   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 [MUSTSHOULD].

1.1 Information Classification Policies

   Information is an asset, but not all information has the same value
   for a business. Not all information needs to be protected as
   strongly as other information.

   Research and development plans, marketing strategies and
   manufacturing quality specifications developed and used by a company
   provide competitive advantage. This type of information needs

stronger protective measures than other information, which if
disclosed or modified, would cause moderate to severe damage to the
company.

Other types of information such as internal organization charts,
employee lists and policies may need little or no protective measures
based on value the organization places on it.

A corporate information classification policy defines how its
information assets are to be protected. It provides guidance to
employees on how to classify information assets. It defines how to
label and protect an asset based on its classification and state
(e.g., facsimile, electronic transfer, storage, shipping, etc.).

1.2 Access Control and Security Labels

"Access control" is a means of enforcing authorizations. There are a
variety of access control methods that are based on different types
of policies and rely on different security mechanisms.

   - Rule based access control is based on policies that can be
     algorithmically expressed.

   - Identity based access control is based on a policy which applies
     explicitly to an individual person or host entity, or to a defined
     group of such entities. Once identity has been authenticated, if
     the identity is verified to be on the access list, then access is
     granted.

   - Rank base access control is based on a policy of hierarchical
     positions in an organization. It is based on who you are in the
     company structure. A rank-based policy would define what
     information that the position of Partner or Senior Consultant could
     access.

   - Role based access control is based on a policy of roles in an
     organization. It may or may not be hierarchical. It is based on
     who you are in the company. The role-based policy would define
     what information that the role of Database Administrator, Network
     Administrator, Mailroom Clerk or Purchaser could access.

Rule, rank and role-based access control methods can rely on a
security label as the security mechanism to convey the sensitivity or
classification of the information. When processing an S/MIME
encapsulated message, the sensitivity information in the message's
security label can be compared with the recipient's authorizations to
determine if the recipient is allowed to access the protected
content.

An S/MIME security label may be included as a signed attribute in the
inner (or only) signature or the outer signature. In the case of a
triple-wrapped message as defined in RFC 2634, the inner signature
would be used for access control decisions related to the plaintext
original content, while the outer signature would be used for access
control decisions related to the encrypted message.

1.3 User Authorizations

   Users need to be granted authorizations to access information that
   has been classified by an authority. The sending and receiving
   agents need to be able to securely determine the user's
   authorizations for access control processing.

   X.509 [X.509] and the Internet profile for X.509 certificates
   [CERTCRL] do not define the means to represent and convey
   authorizations in a certificate.

   X.501 [X.501] defines how to represent authorization in the form of a
   clearance attribute. The clearance attribute identifies the security
   policy in force to which a list of possible classifications and
   security categories relates.

   X.501 also notes two means for binding the clearance to a named
   entity: an Attribute Certificate and a Certificate extension field
   (e.g., within the subjectDirectoryAttribute extension).

   RFC 3281 [AC509] defines a profile of X.509 Attribute Certificate
   (AC) suitable for use with authorization information within Internet
   Protocols. One of the defined attributes is Clearance, which carries
   clearance (security labeling) information about the AC owner. The
   syntax for Clearance is imported from X.501.

2. Developed Examples

2.1 Classification Policies

   The following describes the information classification policies in
   effect at 3 companies.

2.1.1 Amoco Corporation

   The description for the Amoco information classification policy was
   taken from the Amoco Computer Security Guidelines. Amoco classifies
   its information assets based on confidentiality and integrity and
   defines 3 hierarchical classifications for each. The confidentiality

and integrity polices are independent, so either or both may be
applied to the information. Amoco also defines an availability
classification for time critical information.

HIGHLY CONFIDENTIAL - Information whose unauthorized disclosure will
cause the company severe financial, legal or reputation damage.
Examples: Certain acquisitions, bid economics, negotiation
strategies.

CONFIDENTIAL - Information whose unauthorized disclosure may cause
the company financial, legal, or reputation damage. Examples:
Employee Personnel & Payroll Files, some interpreted Exploration
Data.

GENERAL - Information that, because of its personal, technical, or
business sensitivity is restricted for use within the company.
Unless otherwise classified, all information within Amoco is in this
category.

MAXIMUM - Information whose unauthorized modification and destruction
will cause the company severe financial, legal, or reputation damage.

MEDIUM - Information whose unauthorized modification and destruction
may cause the company financial, legal, or reputation damage.
Examples: Electronic Funds, Transfer, Payroll, and Commercial Checks.

MINIMUM - Although an error in this data would be of minimal
consequence, this is still important company information and
therefore will require some minimal controls to ensure a minimal
level of assurance that the integrity of the data is maintained.
This applies to all data that is not placed in one of the above
classifications. Examples: Lease Production Data, Expense Data,
Financial Data, and Exploration Data.

CRITICAL - It is important to assess the availability requirements of
data, applications and systems. A business decision will be required
to determine the length of unavailability that can be tolerated prior
to expending additional resources to ensure the information
availability that is required. Information should be labeled
"CRITICAL" if it is determined that special procedures should be used
to ensure its availability.

2.1.2 Caterpillar, Inc.

The description for the Caterpillar information classification policy
is taken from the Caterpillar Information Protection Guidelines.
Caterpillar classifies its information assets based on
confidentiality and defines 4 hierarchical classifications.

Caterpillar Confidential Red - Provides a significant competitive
advantage. Disclosure would cause severe damage to operations.
Relates to or describes a long-term strategy or critical business
plans. Disclosure would cause regulatory or contractual liability.
Disclosure would cause severe damage to our reputation or the public
image. Disclosure would cause a severe loss of market share or the
ability to be first to market. Disclosure would cause a loss of an
important customer, shareholder, or business partner. Disclosure
would cause a long-term or severe drop in stock value. Strong
likelihood somebody is seeking to acquire this information.

Caterpillar Confidential Yellow - Provides a competitive advantage.
Disclosure could cause moderate damage to the company or an
individual. Relates to or describes an important part of the
operational direction of the company over time. Important technical
or financial aspects of a product line or a business unit.
Disclosure could cause a loss of Customer or Shareholder confidence.
Disclosure could cause a temporary drop in stock value. A likelihood
that somebody could seek to acquire this information.

Caterpillar Confidential Green - Might provide a business advantage
over those who do not have access to the same information. Might be
useful to a competitor. Not easily identifiable by inspection of a
product. Not generally known outside the company or available from
public sources. Generally available internally. Little competitive
interest.

Caterpillar Public - Would not provide a business or competitive
advantage. Routinely made available to interested members of the
General Public. Little or no competitive interest.

2.1.3 Whirlpool Corporation

The description for the Whirlpool information classification policy
is taken from the Whirlpool Information Protection Policy. Whirlpool
classifies its information assets based on confidentiality and
defines 3 hierarchical classifications. The policy states that:

"All information generated by or for Whirlpool, in whatever form,
written, verbal, or electronic, is to be treated as WHIRLPOOL
INTERNAL or WHIRLPOOL CONFIDENTIAL. Classification of information in
either category depends on its value, the impact of unauthorized
disclosure, legal requirements, and the manner in which it needs to
be used by the company. Some WHIRLPOOL INTERNAL information may be
authorized for public release."

WHIRLPOOL CONFIDENTIAL - A subset of Whirlpool Internal information, the unauthorized disclosure or compromise of which would likely have an adverse impact on the company's competitive position, tarnish its reputation, or embarrass an individual. Examples: Customer, financial, pricing, or personnel data; merger/acquisition, product, or marketing plans; new product designs, proprietary processes and systems.

WHIRLPOOL INTERNAL - All forms of proprietary information originated or owned by Whirlpool, or entrusted to it by others. Examples: Organization charts, policies, procedures, phone directories, some types of training materials.

WHIRLPOOL PUBLIC - Information officially released by Whirlpool for widespread public disclosure. Example: Press releases, public marketing materials, employment advertising, annual reports, product brochures, the public web site, etc.

The policy also states that privacy markings are allowable. Specifically:

For WHIRLPOOL INTERNAL, additional markings or caveats are optional at the discretion of the information owner.

For WHIRLPOOL CONFIDENTIAL, add additional marking or caveats as necessary to comply with regulatory or heightened security requirements. Examples: MAKE NO COPIES, THIRD PARTY CONFIDENTIAL, ATTORNEY-CLIENT PRIVILEGED DOCUMENT, DISTRIBUTION LIMITED TO _____, COVERED BY A NON-ANALYSIS AGREEMENT.

2.2 S/MIME Classification Label Organizational Examples

RFC 2634 [ESS] defines the ESSSecurityLabel syntax and processing rules. This section builds upon those definitions to define detailed example policies.

2.2.1 Security Label Components

The examples are detailed using the various components of the eSSSecurityLabel syntax.

2.2.1.1 Security Policy Identifier

A security policy is a set of criteria for the provision of security services. The eSSSecurityLabel security-policy-identifier is used to identify the security policy in force to which the security label relates. It indicates the semantics of the other security label components.

   For the example policies, the following security policy object
   identifiers are defined:

   -- S/MIME Working Group Object Identifier Registry
   id-smime OBJECT IDENTIFIER ::= { iso(1) member-body(2) us(840)
                                    rsadsi(113549) pkcs(1) pkcs-9(9) 16 }

   -- S/MIME Test Security Policy Arc
   id-tsp OBJECT IDENTIFIER ::= { id-smime 7 }

   -- Test Security Policies
   id-tsp-TEST-Amoco   OBJECT IDENTIFIER ::= { id-tsp 1 }
   id-tsp-TEST-Caterpillar   OBJECT IDENTIFIER ::= { id-tsp 2 }
   id-tsp-TEST-Whirlpool   OBJECT IDENTIFIER ::= { id-tsp 3 }

2.2.1.2 Security Classification

   The security classification values and meanings are defined by the
   governing company policies. The security-classification values
   defined are hierarchical and do not use integers 0 through 5.

   Amoco-SecurityClassification ::= INTEGER {
     amoco-general (6),
     amoco-confidential (7),
     amoco-highly-confidential (8) }

   Caterpillar-SecurityClassification ::= INTEGER {
     caterpillar-public (6),
     caterpillar-green (7),
     caterpillar-yellow (8),
     caterpillar-red (9) }

   Whirlpool-SecurityClassification ::= INTEGER {
     whirlpool-public (6),
     whirlpool-internal (7),
     whirlpool-confidential (8) }

2.2.1.3 Privacy Mark

   Privacy marks are specified the Whirlpool policy. The policy
   provides examples of possible markings but others can be defined by
   users as necessary (though no guidance is given). The Whirlpool
   policy provides the following examples: MAKE NO COPIES, THIRD PARTY
   CONFIDENTIAL, ATTORNEY-CLIENT PRIVILEGED DOCUMENT, DISTRIBUTION
   LIMITED TO _____, and COVERED BY A NON-ANALYSIS AGREEMENT.

The Amoco policy does not identify any privacy marks but the
classification labels defined for availability and integrity would be
most appropriately displayed here. The CRITICAL, MAXIMUM, MEDIUM,
and MINIMUM labels are examples of information classifications that
are not used for access control.

In general, the privacy marks should provide brief but clear
direction to the user on how to handle the information.

2.2.1.4 Security Categories

Security categories or caveats are not specified in any of the sample
policies. However, they are used in at least 2 of the companies.
Though the security categories are not defined formally in their
security policies, once locally defined they are formal and are to be
enforced. The security categories are defined when necessary to
provide identifiable proprietary information more granular access
control. A category can be based organizationally or by project
(i.e., Legal Only or Project Vallor).

2.2.1.4.1 Syntax

Security categories are represented in the RFC 2634 ESSSecurityLabel
(to specify the sensitivity of labeled data) and X.501 Clearance
attribute (to specify an entity's authorizations) using the following
syntax.

```
SecurityCategories ::= SET SIZE (1..ub-security-categories)
                       OF SecurityCategory

ub-security-categories INTEGER ::= 64

SecurityCategory ::= SEQUENCE {
  type  [0] OBJECT IDENTIFIER
  value [1] ANY DEFINED BY type } -- defined by type
```

One example of a SecurityCategory syntax is SecurityCategoryValues,
as follows.

When id-securityCategoryValues is present in the SecurityCategory
type field, then the SecurityCategory value field could take the form
of:

```
SecurityCategoryValues ::= SEQUENCE OF UTF8String
```

2.2.1.4.2 Use

   An organization will define a securityCategoryType OID representing
   the syntax for representing a security category value within their
   security policy.

   For the example security category syntax, a UTF8String is used to
   convey the security category value that applies to the labeled
   message. Access MUST be restricted to only those entities who are
   authorized to access every SecurityCategoryValue. Access is
   authorized if the ESSSecurityLabel SecurityCategoryValue EXACTLY
   matches the Clearance SecurityCategoryValue.

2.2.2 Attribute Owner Clearance

   The security clearance and category authorizations for the user are
   defined in the clearance attribute.

2.2.2.1 Amoco User

   Clearance:
     policyId: 1 2 840 113549 1 9 16 7 1
     classList: amoco-general   (6),
                 amoco-confidential   (7),
                 amoco-highly-confidential (8)

2.2.2.2 Caterpillar User

   Clearance:
     policyId: 1 2 840 113549 1 9 16 7 2
     classList: caterpillar-public   (6),
                 caterpillar-confidential-green (7),
                 caterpillar-confidential-yellow (8),
                 caterpillar-confidential-red   (9)

2.2.2.3 Whirlpool User

   Clearance:
     policyId: 1 2 840 113549 1 9 16 7 3
     classList: whirlpool-public   (6),
                 whirlpool-internal   (7),
                 whirlpool-confidential (8)

2.2.3 Security Category Example

   This section includes an example RFC 2634 ESSSecurityLabel including
   the example Security Category syntax. This section also includes
   example X.501 Clearance attributes. One of the example Clearance
   attributes includes a set of authorizations that pass the access
   control check for the example ESSSecurityLabel. The other example
   Clearance attributes each include a set of authorizations that fail
   the access control check for the example ESSSecurityLabel.

   These examples use the id-tsp-TEST-Whirlpool OID defined in section
   2.2.1.1. Assume that the security policy identified by id-tsp-TEST-
   Whirlpool defines one securityCategoryType OIDs as follows:

   id-tsp-TEST-Whirlpool-Categories OBJECT IDENTIFIER ::= { id-tsp 4 }

   Example ESSSecurityLabel:
    security-policy-identifier: id-tsp-3
    security-classification: 8
    privacy-mark: ATTORNEY-CLIENT PRIVILEGED INFORMATION
    security-categories: SEQUENCE OF SecurityCategory

   SecurityCategory #1
     type: id-tsp-4
     value: LAW DEPARTMENT USE ONLY

   Example Clearance Attribute #1 (passes access control check):

   Clearance:
     policyId: id-tsp-3
     classList BIT STRING: Bits 6, 7, 8 are set to TRUE
     securityCategories: SEQUENCE OF SecurityCategory

   SecurityCategory #1
     type: id-tsp-4
     value: LAW DEPARTMENT USE ONLY

   Example Clearance Attribute #2 (fails access control check because
   SecurityCategoryValues do not match):

   Clearance:
     policyId: id-tsp-3
     classList BIT STRING: Bits 6, 7, 8 are set to TRUE
     securityCategories: SEQUENCE OF SecurityCategory

```
   SecurityCategory #1:
     type:  id-tsp-4
     value:  HUMAN RESOURCES USE ONLY
```

2.2.4 Additional ESSSecurityLabel Processing Guidance

   An implementation issue can be the mapping of the security label
   values to displayable characters. This is an issue for users who
   want to develop and retire their own classifications and categories
   on a regular basis and when the values are encoded in non-human
   readable form. Applications should provide a means for the
   enterprise to manage these changes. The practice of hard coding the
   mapping into the applications is discouraged.

   This issue is viewed as local issue for the application vendor, as
   the solution does not need to be interoperable between vendors.

   An approach is the use of a Security Policy Information File (SPIF)
   [ISO15816]. A SPIF is a construct that conveys domain-specific
   security policy information. It is a signed object to protect it
   from unauthorized changes and to authenticate the source of the
   policy information. It contains critical display information such as
   the text string for security classifications and security categories
   to be displayed to the user, as well as additional security policy
   information.

   Another implementation issue can be obtaining the recipient's
   certificate when sending a signed-only message with a security label.
   Normally the recipient's certificate is only needed when sending an
   encrypted message. Applications will need to be able to retrieve the
   recipient's certificate so that the recipient's clearance information
   is available for the access control check.

3. Security Considerations

   All security considerations from RFC 2630 [CMS] and RFC 2634 [ESS]
   apply to applications that use procedures described in this document.

References

    [AC509]   Farrell, S. and R. Housley, "An Internet Attribute
                   Certificate Profile for Authorization", RFC 3281, April
                   2002.

    [CERTCRL]   Housley, R., Polk, W., Ford, W. and D. Solo, "Internet
                   X.509 Public Key Infrastructure Certificate and
                   Certificate Revocation List (CRL) Profile", RFC 3280,
                   April 2002.

    [CMS]   Housley, R., "Cryptographic Message Syntax", RFC 2630,
                   June 1999.

    [ESS]   Hoffman, P., Editor, "Enhanced Security Services for
                   S/MIME", RFC 2634, June 1999.

    [MUSTSHOULD]   Bradner, S., "Key Words for Use in RFCs to Indicate
                   Requirement Levels", BCP 14, RFC 2119, March 1997.

    [X.501]   "ITU-T Recommendation X.501: Information Technology –
                   Open Systems Interconnection - The Directory: Models",
                   1993.

    [X.509]   "ITU-T Recommendation X.509 (1997 E): Information
                   Technology – Open Systems Interconnection – The
                   Directory: Authentication Framework", June 1997.

    [ISO15816]   "Information Technology - Security Techniques - Security
                   Information Objects for Access Control", ISO/IEC FDIS
                   15816:2000.

Author's Address

     Weston Nicolls
     Forsythe Solutions
     7500 Frontage Rd
     Skokie, IL 60077

     Phone: (847) 763-2370
     EMail: wnicolls@forsythesolutions.com

Full Copyright Statement

Acknowledgement