INFORMATION LABELING - CROSS-DOMAIN SOLUTIONS

De heren ir. B.J. te Paske, ir. D. Boonstra, ir. H. Hut en ir. H.A. Schotanus (allen TNO)

De auteurs van dit artikel, de heren ir. B.J. te Paske, ir. D. Boonstra, ir. H. Hut en ir. H.A. Schotanus, werken bij TNO-ICT afdeling Security, een afdeling met zo'n 25 security professionals. De auteurs werken aan diverse security en informatiebeveiligingsprojecten voor zowel Defensie als voor andere markten. In dit artikel wordt nader ingegaan op het conceptuele model van Cross-Domain Solutions en wordt Information Labeling nader uit de doeken gedaan. Dit artikel is tot stand gekomen binnen het Informatiebeveiligingsonderzoeksprogramma voor Defensie.

INTRODUCTION

The importance of sharing information in networked military operations is commonly recognized. An important concept in this context is Network-Enabled Capabilities (NEC). Better integrated networks mean that sharing relevant military information will be easier and quicker and that more people can be reached than before. The ultimate goal is an integrated and coordinated deployment of all capabilities. In the NEC vision everyone and everything is connected in order to allow for information sharing: 'the right information, at the right place, at the right time - but not too much information'. This requires harmonization and change; harmonization of the existing capabilities and change in terms of doctrine, processes, personnel, culture and organization.

The technical basis for networked operations lies in a secure, robust and extensive federation of networks, a large network consisting of a number of smaller individual networks. This future network is known by the term Network and Information Infrastructure (NII), which is in fact an ever-developing coalition-wide multinational military 'intranet' alike network. This network can be considered a collection of different security domains.

Several challenges exist in making it possible to share information between security domains. Some of these are addressed by the NATO RTO study on 'XML in Cross Domain Security Solutions' (IST-068). The aim of IST-068 is defined as 'improving the possibility to share information between national security domains, by using XML security solutions to properly describe the security properties of information objects'. This reflects two important notions:

1. Sharing information between security domains presumes a way of determining whether a specific information object may or may not be shared.
2. Each information object has a set of security properties that are relevant in this process.

Such a set of security properties may be referred to as a '(security) label'. It should be noted that the relevance of this concept is not limited to digital information objects. On the contrary, long before the digital era our Defense organizations were used to working with security classifications on paper documents. An important question is if these traditional concepts and practices are equally valid in the present digitalized world, and if differences exist that incur new challenges or possibilities. One of these possibilities is to deploy the eXtensible Markup Language (XML) to define the content, structure and format of a security label in a standardized way.

The goal of this document is to create and present a Cross-Domain Solution conceptual model which RTG-031 can use to explain some properties of Cross-Domain security issues in terms of requirements, example scenarios, security perspectives, algorithms, access control decision functions and so on.

PAPER-BASED INFORMATION OBJECTS: A PERSPECTIVE

As explained before, the basis for our conceptual model lies in the traditional way of working with classified information and the reasons for classifying information in the first place. Within Defense organizations there is a high use of classified information. Classified information is sensitive information to which access is restricted by law or regulation to particular classes of people and classes of environments and which is needed for the correct performance of people's duties.

Information classifications are made explicit by marking documents with an exclusivity marking indicating among others its level of sensitivity. A marking is human understandable, formatted text that is applied to printed or displayed information. A marking also represents restrictions on the handling or distribution of that information as required by a security policy. Exclusivity markings are a means for controlling information flows. The ability to control information flows is a precondition for effective information security.

An exclusivity marking may consist of three elements:

- A policy identifier
- A classification
- A category

The policy identifier defines which authority is responsible for the exclusivity marking as well as the regulations to work with classified information. Examples are 'NATO', 'STG' and 'WEU'.

The classification indicates the sensitivity level of an information object. Typical examples are 'restricted', 'confidential', 'secret' and 'top secret'.

A category is an additional indication of sensitivity, application or treatment that may either limit or broaden the access to a piece of information. Examples are 'releasable to WEU', 'SFOR only' or 'Releasable for Internet Transmission'.

Typically the policy and classification are well defined while the categories are more free-form.

In terms of these definitions, the use of security labels or exclusivity markers serves the following purposes:

1. Restricting access to information objects within a security domain. The category field in the exclusivity marker may be used to limit access to a 'Community-of-Interest' with a need-to-know;
2. Preventing information objects from being shared with anybody outside the security domain. The policy and classification fields may be inspected to ensure that for example information objects like documents with a national classification may not be shared with other nations or NATO.
3. Controlling the use of information objects over different security domains, based on a shared security policy. E.g. mission-specific classifications may be used

to share information between allies in a controlled way.

Important to note is that for the first and second purpose, the security label is only locally used within one security domain. This means that no inter-domain agreements are required on the format of the label and the policy associated with it. Only when security labels are used between security domains, it must be clear for all parties involved what the syntax and semantics of a label may be and how it must be interpreted.

To conclude this summary of how information classification is used in the 'paper world', we list a number of measures that are commonly used to address threats against the confidentiality, integrity and availability of sensitive information objects:

-Registration of information objects, providing them with unique identifiers. With a proper registration it is possible to administrate where documents are located in the domain.

-Logging of meta-data including name of the author who created the information object, name of the person who approved the classification and the date and time of the approval.

-Storing classified information in a secure way.

-Formal security clearance required for handling classified information. The clearance process involves a background investigation which is more profound if the clearance allows working with higher exclusivity markings.

-Definition of formal policies for sharing information within the security domain and between security domains.

DIGITAL INFORMATION IN A NETWORKED ENVIRONMENT

In the 'paper world' information has a physical form factor and must be kept in a physical location (e.g. a safe in a room in a building) and shared through manual handover. In the digital era, we have grown used to working with digital information which is stored and shared across a network infrastructure. Although many of the concepts and threats related to information protection are equally valid in the digital era, it has introduced additional issues to be taken into account.

In general, it can be stated that digital information, compared to physical documents:

- can be more easily shared with, or accessed by, others; in addition, access or sharing of information frequently means duplication of that information.
- can be more easily modified; furthermore, it is easier to modify information in such a way that modifications remain unnoticed.
- can take many forms; many different document formats exist, not to mention multimedia, real-time streams and so on.

The first aspect may introduce new threats related to confidentiality, while the second may bring new threats related to integrity. The third aspect makes it more challenging to implement uniform mechanisms for controlling information flows (which, as we discussed earlier, is a precondition for achieving information security).

While this implies that Information Technology (IT) and digitalization have confronted us with a range of challenges, IT may also offer us the means to implement appropriate measures. A number of functional components may be implemented, as depicted in Figure 2.

The functional components in a cross-domain information exchange scenario include:

-Information Exchange Gateway (IEG)

An IEG aims to protect a security domain against certain threats. It is typically implemented as a demilitarized zone (DMZ) which hides the protected network, monitors traffic and provides a network interface for certain services. As such an IEG provides protection at the network infrastructure level.

-Labeling

Labeling involves the application and management of digital security labels in relation to information objects. These security labels must be meaningful within the context of a security policy.

-Release mechanism

A release mechanism determines whether or not an information object may be released from a security domain. This decision may be based on inspection of the security label, and following the applicable security policy. A release mechanism, as opposed to an IEG, provides control at the information level. The release mechanism is an enforcement point of the security policy with respect to information that will leave the security domain.

-Acceptance mechanism

An acceptance mechanism controls which information is accepted into a network from other security domains. This decision may be based on an information label and should be in accordance with the applicable security policy. The acceptance mechanism is the reverse of the release mechanism; it is an enforcement point of the security policy with respect to what information is permitted to enter the security domain.

-Transference mechanism

A transference mechanism is used to transfer the document with a set of properties pertaining to that document – for example in the form of a label – from one domain to another. The participating domains must be in agreement about which policies apply and the format and content of the set of properties. This is in part a technical mechanism, but for a large part this consist of arrangements about the policies and the meaning of the properties.

It should be noted that the case of information exchange between NATO and some NATO nation is just one example where these functional components may be implemented. They are equally relevant when information exchange takes place between two system-high security domains, e.g. between two security domains within one nation.

In this document the following concepts are used to discuss the use of security markings and labeling:

Information object
An information object is any piece of information that, as a whole, represents a certain meaning within a specific context. An information object may use different formats and document lay-outs.

Security label
A security label is a set of security properties related to the information object. These security properties can be used to determine among others if an information object can be released to a particular user or domain.

Security policy
In the context of information exchange, a security policy is a set of rules for protecting information against unauthorized disclosure or modification. An implementation of the security policy includes defining how those rules apply to a particular information object.

Security domain
We follow the NIST definition [Reference: http://csrc.nist.gov/publications/nistir/NISTIR-7298_Glossary_Key_Infor_Security_Terms.pdf] of a security domain as 'a collection of entities to which applies a single security policy executed by a single authority'. Each security domain is responsible for establishing and maintaining its own level of information security while all security domains together are responsible for establishing and maintaining the overall level of information security.

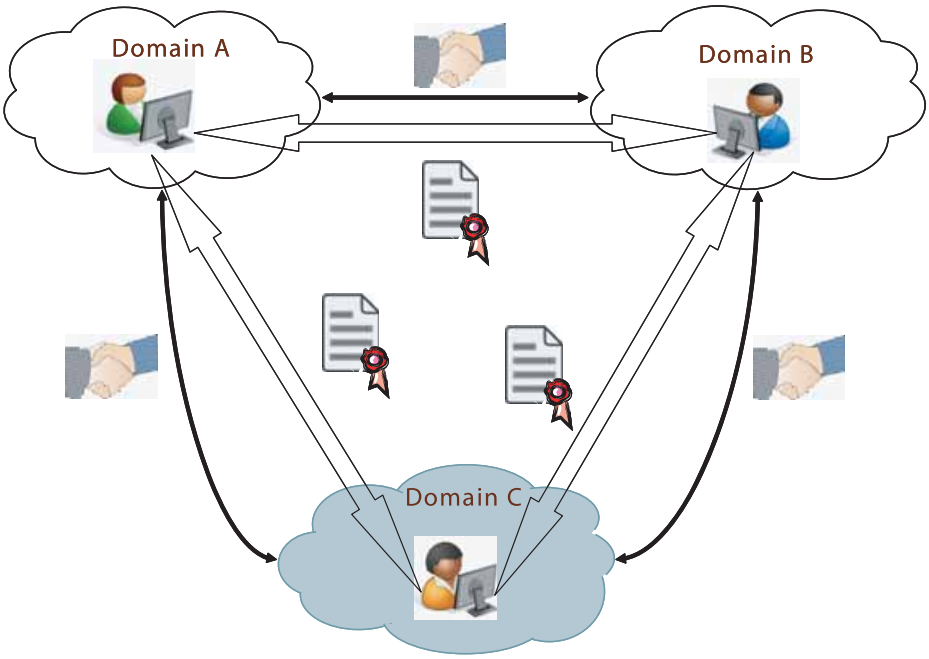

Figure 1: Labeling across security domains

In general, when information flows from one security domain to another, there are three different types of boundaries that information objects must face:
1. The organizational boundary
2. The marking (or policy) boundary
3. The classification level boundary

When the organizational boundary is crossed, literally the information is handed over to a different organization. The responsibility for correctly handling the information is therefore also directly handed over to that organization. An example of such a boundary is a connection between a NATO secret network operated by NATO and a NATO secret network operated by a NATO nation.
The second boundary is the marking or policy boundary. This happens for example when a national network with national classified information is connected to a national network with NATO classified material of an equivalent level within the security domain of one organization. Imagine for example a national classified network connected to a NATO classified network both operated and maintained by a NATO nation.

The third boundary is the classification level boundary. Imagine for example a NATO Secret network connected to a NATO Confidential network.

Note that in all three cases classified information objects may need to be reclassified before information can be exchanged; that is the security marking might need to be changed. Also note that the category that may be included in the exclusivity marking is not explicitly included in this model. The reason is that it is hard to formalize as a separate entity in this model but still it must be taken into account because it might allow or disallow certain information flows.

In many cases however the information will be required to cross more than one of these boundaries. For example when a NATO nation connects its national secret network to a NATO Secret network operated by NATO. In that case both the organization and the marking changes. This implies that there might be a need to apply different security mechanisms to such connections. Hence a strict separation in these three boundaries is required when discussing connections between different security domains.

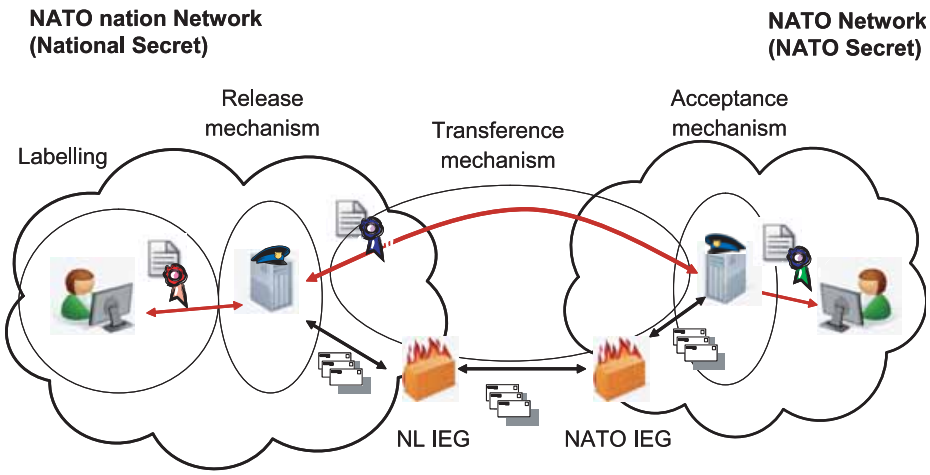

Figure 2: Functional components in cross-domain information exchange

LABELING ISSUES

The conceptual model displayed in Figure 2 and Figure 3 forms a basis for the realization of a networked operational environment. There are nevertheless a large number of issues that have to be filled in. This section attempts to divide the problem of labeling and discusses these issues separately.

1. What information should be contained in a label?
There is a plethora of meta-information about the information object that can be stored in the label itself. It must be determined which information is required to be included in the label. It is evident that the label must contain the exclusivity markings pertaining to the document. Additionally a date of the classification, the individual that assessed the classification and the classification validity period can typically be found on paper information. It seems therefore logical that this information is also included in the digital security label. There are also a number of other elements that might be included in the label, such as an identifier of a community of interest and the title, author, date and version of the document. Nevertheless it is perhaps not necessary to determine all the fields beforehand, but develop a label in such a manner that it can easily be extended.

2. Is the label contained in the document or is it a separate meta-object?
There are two ways to look at a label; one is that it is purely used to indicate the classification and as such can be included in the information itself. The second is to have a separate meta-object that contains the classification as well as other meta-information about the classification. The first one has the advantage that the label is never separated from the document, which makes it much easier to retrieve the label, given the document or vice versa. This however yields the problem that existing applications have to be able to work with this label inside the document. This can be partly alleviated by creating a new object that contains both the label and the document. Yet this also requires a means for applications to extract the document from the container object. A separate meta-object that consists of only the label does not suffer from any of these issu-

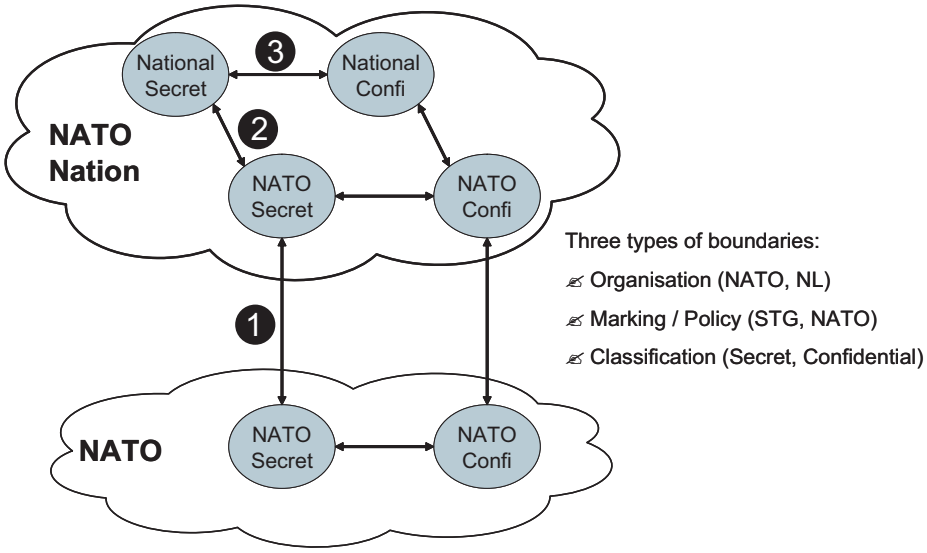

Figure 3: Three types of boundaries

es, as it is application independent. But this results in other issues, including retrieving a label for a document which needs a separate solution.

3. How can we bind a label object to an information object?
The label and document are two separate entities that also may be stored independently; however there must be a 1-to-1 relation between the label and the information object. Thus a label must only be valid for a single information object and a document should only have a single label. Especially the second issue can be problematic when there are multiple labels with conflicting information. Such a case is not altogether unlikely. For example when two individuals separately decide to label a document at more or less the same time, or when a labeled document is copied to another domain and the label is stripped away in that domain. Another individual may create a new label for the document. It is therefore necessary to either prevent or detect it, which may be found technically difficult, particularly when labels can exist in other domains. Or in the least it must be possible to resolve any conflicts in labels.

4. How can we secure the binding of information object to the label?
The relation between the label and the information object should be a reliable and secure mechanism as this relation is the core of the purpose for labeling. It must be possible to ascertain that the label is bound to a single entity of information and that the contents of the label or the document cannot be changed without invalidating the label.

5. How can the semantic meaning of a label be determined?
The label refers to a specific policy. The policy describes how the document with said label must be treated within the particular security domain. This means that it is necessary to have agreements in place between the parties involved in the information exchange on how the information will be treated.

6. How can we inform an end-user about the used policies?
With just the label and the policy in place, we are not quite there yet. There is also a necessity to disseminate the policy associated with the label to the end-user and ensure that the user adheres to the policy. The user must be made aware of both the label itself and the policy.
The main problem here is that with more and more connections to other domains, this will open the door to many more different policies. Consider for example different policies associated with missions, with the NATO organization, and with NATO or non-NATO nations. The differences can vary from minor adjustments to a completely different policy. In any case the user must be aware of which policy to apply to a specific document. In addition the user must also be aware of the specifics of the policy involved. This will have to be dealt with in a higher level and requires management of all the policies.

7. How can the label be visualized or presented in the user interface?
The end-user must somehow be able to work with labels. Therefore the label and the mechanism must be integrated in the user environment, including the applications, so that a user is aware of the label associated with the information. This implies that an application involved must be capable of working with a labeling mechanism to extract the information from the label which it needs to present that to the end-user. Applications – especially COTS products – will not likely be able to do this out of the box.

8. How can the user securely create a label for a given object?
One of the most difficult tasks is how the label can be created and securely bound to a specific document. This means that the user must have proper authorization to do so, which must be validated. And also the user must be able to ascertain that the label is bound to the information that the user intended. This is mostly a technical issue and requires a trusted path between the labeling mechanism and the information storage. The main threat here is that during or before the labeling process the information is altered without the user's knowledge. There are many ways this can be achieved; for example the document could be overwritten when it is stored on a shared disc right after the user saves it but right before the information is read by the labeling mechanism (race condition). A second more insidious example is malicious software or otherwise malfunctioning software that can slip information into an existing document. Thus within the application the information that is to be labeled must be frozen – unalterable – when the user indicates that it is going to be labeled. The amount of certainty required depends largely on the purpose of the labeling mechanism and the applicable security level.

9. What is needed to validate and enforce the label?
Creating a label for a given information object is only one side of the medallion. It is also needed to develop methods to validate labels and make decisions about labels. It may even be needed to develop a means to translate or associate a label from one domain to another. What specifically is needed however largely depends on the way the label is used. If a label for example is only used to indicate the classification of an information object to a user, the user must be able to verify the authenticity of the label. But if the label is also used to enforce a policy – such as 'only documents classified NATO Unclassified may be shared with security domain XYZ' – an enforcement point – a guard - is also needed. In the latter case, this will put much higher assurance requirements on the labeling infrastructure.

BRONNEN:

Voor meer informatie kunt u contact opnemen met:
Ir. A.C.M. Smulders
Programmamanager V813 programma informatiebeveiliging
andre.smulders@tno.nl
Tel: +31 15 2857272