

Protection Profiles for Discretionary Information Flow Control (DIC)

Dr. Steffen Lange

Dr. Andreas Nonnengart

Christian Stühle

Roland Vogt

Marcel Weinand

Deutsches Forschungszentrum für
Künstliche Intelligenz (DFKI GmbH)
[German Research Center for Artificial Intelligence]
Stuhlsatzenhausweg 3
D-66123 Saarbrücken
GERMANY

Bundesamt für Sicherheit in der
Informationstechnik (BSI)
[German Information Security Agency]
Godesberger Allee 185-189
D-53175 Bonn
GERMANY

Phone: +49-681-302-4131

Email: dic-dfki@DFKI.de

Phone: +49-1888-9582-152

Email: weinand@bsi.bund.de

Abstract

Today's information flow control in operating systems is usually based on access control functions, encryption procedures and firewalls. Additionally, the transport of information in open networks can be secured by Virtual Private Networks (VPNs) or by the use of Public Key Infrastructures (PKIs). However, as it turns out, such security mechanisms are difficult to use and require an enormous administration effort. The consequence is that information that actually would need protection, often flows unprotected across public networks.

On behalf of the German Information Security Agency (Bundesamt für Sicherheit in der Informationstechnik – BSI) the German Research Center for Artificial Intelligence (Deutsches Forschungszentrum für Künstliche Intelligenz – DFKI GmbH) is going to create two closely related Common Criteria Protection Profiles for Discretionary Information Flow Control (DIC). The project, which is driven by the German federal and state data protection commissioners, is meant to help for better Privacy Enhancing Technologies (PETs). Its task is the specification of requirements for the control of all information flows within computers in accordance with adjustable rules. Such rules (like e.g. encrypted storage, encrypted transmission, signed transmission) are to be provided by the user who is responsible for the security of the information. In general they can be derived from given technical, organizational or legal regulations.

The intended security functionality should be transparent for users as well as for applications. Therefore, the concept of the Protection Profiles proceeds from the creation of a security layer which resides below the application layer and is integrated in the operating system.

Protection Profiles for Discretionary Information Flow Control (DIC)

Introduction

The task of the protection profiles is to provide requirements for the protection of information flows of an IT-system transparently for the users. That is why the TOE controls the admission of an information flow according to definable information flow rules. The TOE supports in particular IT users with little competence in IT issues in the enforcement of protection of information, as far as the aspects confidentiality, integrity and/or authenticity are concerned. The TOE security features can be considered as a useful addition to well-established security concepts such as access control, transmission protection, firewalls or Virtual Private Networks. The TOE can be used in different areas like:

- E-commerce (data warehouses etc.),
- E-government (tendering procedures, application procedures etc.),
- Health System (electronic patient records etc.) and
- Tele and media services (telework etc.).

A combination of security mechanisms corresponding to the security requirements of an information flow can be allocated to every single information flow. For the controlled information these mechanisms guarantee selectively the protection of the

- integrity by electronic signature,
- confidentiality by encrypting, as well as
- authenticity by electronic certificates.

In this context the maintenance of confidentiality serves to avoid undesirable notice of locally stored user data (e.g. when data carriers are stolen or the information processing is inappropriate) and of user data during the transmission of messages. Integrity and authenticity are particularly important when it comes to commercial transactions (e.g. electronic orders and electronic payments).

Another protection mechanism is the restriction of information processing to certain subjects (e.g. applications). This supports to realise in a technical sense the appropriateness of information processing in accordance with data protection regulations.

The TOE operates almost completely transparent for the concerned subjects (e.g. applications) and for users of the IT-system. An adjustment of the applications employed in the IT-system only has to be done as far as the TOE has to be provided with information about the respective information flow in order to maintain its control. Flexible configuration options enable an individual and consistent adjustment of the TOE to the protection requirements of the operator of the IT-system.

The protection profile abstracts away from technical details of the demands on the TOE so that a realisation for a number of various IT-environments is possible such as

- (single and multi user) operating systems,
- data base systems or
- e-mail-clients and e-mail servers.

In this context it should be mentioned that the TOE can be integrated in the respective IT-environment in varied ways. A service process of an operating system such as an e-mail server can be considered as a single subject being controlled by a TOE which is integrated in the operating system. It can also be considered as an independent IT-environment for a TOE which controls clients communicating with the server.

Attacker Model and Security Objectives

The background of the development of the protection profiles is a scenario where three different security needs closely interact. It is characterized by the presence of sensitive information that

1. should be kept in a restricted area,
2. should be protected against unauthorized disclosure resp. manipulation, and
3. should only be processed for serving its purpose.

The main motivation of the presented work comes from the first aspect which is associated with the threat of distributing sensitive information across insecure domains. It is assumed that information flows are explicitly performed by read/write operations and that violations of domain borders are accidentally triggered by authorized users. This means that covert channels as well as fraudulent user behavior are precluded. As a consequence the TOE's information flow control features sufficiently can rely on fairly elementary concepts. Additionally, it should be in the sole discretion of specifically authorized users to override any policy based decision that would otherwise lead to the denial of a certain information flow.

The second aspect is meant to help for the preservation of confidentiality, integrity and/or authenticity of sensitive information. Attacks possibly occur inside secure domains as well as during data transmission. They are originated by unauthorized persons. Note that IT administration personnel is generally attributed as unauthorized attacker. Additionally, in a multi user environment, regular IT users are considered as attackers since they might be interested in disclosing resp. manipulating information of their colleagues. The TOE is required to provide appropriate cryptographic mechanisms in order to prevent unauthorized disclosure and manipulation. Since the protected information is typically stored over a long period, these cryptographic mechanisms have to resist against medium-strength attacks.

Finally, the third aspect is derived from general appropriation regulations. They state that sensitive information should be collected, processed or used only for designated purposes. Therefore the TOE is required to restrict the access to information within the security domain in the sense that information can only be processed by designated subjects.

It is generally assumed that attackers only perform obvious penetration attacks and that the operator of the IT system has a strong interest in the security features of the TOE. Moreover, the TOE administrator is assumed to trustworthily act on behalf of the IT operator while adjusting the otherwise completely separated configuration data. In other words, he never performs any penetration attacks. This implies that there is no need for any kind of control upon the actions of the TOE administrator. Nevertheless, he needs good administration support for avoiding accidental mistakes.

The main security objectives for the TOE are:

1. Provide a discretionary information flow control policy including elementary support for appropriation regulations
2. Prevent unauthorized disclosure resp. unnoticeable manipulation of sensitive information
3. Provide good administration support as well as effective separation of configuration data

Security Policy

This chapter explains the security functional policy (SFP) of the TOE, i.e. the *discretionary information flow control policy*. Here some Common Criteria terms are used which will be explained first.

According to the Common Criteria, describing security functional policies, active units (subjects), passive units (objects), and information are distinguished. Objects may contain information and are the target of operations, carried out by subjects. Security attributes are assigned to Subjects, objects and information from which decisions of SFP's are made.

Concept definition

An **information flow** is the input/output of an IT component from/to any data depository. It is characterised by the indication of a subject (active unit), an object (passive unit) and an operation (read/write).

An information flow is caused by a **subject** which is identified by an active functional unit, such as for example an application (a user program to which tasks on the operation system level can be assigned). The multi user variant [2] of the protection profiles differs from the single user variant [1] in that the user identity is an integral part of the notion of a subject.

The information flowing from/to subjects forms, together with the container where the information is stored, the controlled **object**. Typically, the container is a file; but it can also be considered as a record in a data base system or as a mailbox assigned to an e-mail address.

The **data depository** relates to e.g. a local storage device, an addressee accessible via a network link or a computer address. For a storage device, the data depository can consist of a path information concerning a hierarchical directory. The indication of an e-mail address would be typical for network addressees.

Furthermore two operations are defined:

read(S; I; D) subject S reads information I from object O located at data depository D.

write(S; I; D) subject S writes information I in object O located at data depository D. As the case may be, an information already existing in O is thus deleted.

An **information flow rule** consists of the following components:

- an operation (read/write)
- a set of subjects
- a set of data depositories
- a control flag CF
- an trust flag TF
- a protocol flag PF
- a set of information flow instructions.

The mentioning in the set of the subjects is the condition for the permission of an information flow. With this those subjects can be determined that are meant to be appropriate for the processing of the respective information. To be able to describe a set of subjects in a compact way, it is permitted to use wild cards.

All information flows with any mentioned data depository are carried out in accordance with the specified information flow instructions. To be able to describe a set of data depositories in a compact way, it is permitted to use wild cards.

The control flag CF is set to "True" to indicate that only information flows according to existing information flow rules are permitted to take place with the mentioned data depositories. By doing so, the essential characteristic feature of information flow control can be ensured, namely that information requiring control stays in the controlled domain.

The trust flag TF is set to "True" when the mentioned subjects are permitted to write the information stored at the mentioned data depositories to other data depositories without maintaining the protection. With this it is possible to specify exceptions to the limitation of the controlled domain.

The protocol flag PF is set to "True" when it is necessary to log all requests of information flows which are permitted or denied according to the specific information flow rule.

Via appropriate information flow instructions the authenticity, integrity or confidentiality of the information stored at the mentioned data depositories can be protected. At least, the information flow instructions should be able to express the need for encrypting resp. decrypting information as well as generating resp. verifying electronic signatures on information.

In order to meet different security needs, information flow rules can be collected in a list of information flow rules. We now are considering a specific list of information flow rules and a specific data depository D. An information flow rule R represents the **most specific information flow rule for D** if D is mentioned in R and the list does not contain an information flow rule R* in which D and only some of R's data depositories are mentioned.

With the help of the notion of a most specific information flow rule it is possible to construct a hierarchy of information flow rules and thus to identify those information flow rules that fit best with the current data depository. The example given in Figure 1 depicts three information flow rules and four data depositories. Rule R1 mentions all depositories, rule R2 mentions D2 and D3, and rule R3 mentions D3 and D4. The table associates every data depository with its most specific information flow rules. Note that the sets of data depositories mentioned in rules R2 and R3 overlap. Since in each rule the data depository D3 is mentioned, both R2 and R3 are most specific information flow rules for D3.




Figure 1: Illustration of the notion “most specific information flow rule”

Of course, conflicts as depicted in the preceding example have to be avoided resp. resolved. On the one hand there must be at most one most specific information flow rule for each specific information flow. On the other hand the application of different information flow rules to the same objects must not lead to any loss of information. These considerations lead to a number of appropriate restrictions on the construction of lists of information flow rules. The protection profiles formulate sufficient conditions guaranteeing the unambiguousness of rule selection as well as the consistency of controlled information.

The decisions of the discretionary information flow control policy to be enforced by the TOE are based on the object and subject attributes¹ specified in Table 1.

Category	Attribute	Possible values
Object O	Control Status C(O)	Strong, Weak
Subject S	Security Level L(S)	High, Low

Table 1: Security Attributes

The security attribute **control status** indicates (C(O) = Strong) that information exchange with a controlled object (C(O) = Strong) must take place only according to the specified control rules. Thus, the appropriateness is guaranteed and authenticity, integrity respectively confidentiality of the information kept in this object are protected. The security attribute **security level** indicates whether subject S is allowed (L(S) = Low) to write information in objects with the control status “Weak”.

¹ In the multi user variant subjects are to be considered pairs of user identity and other data needed for the description of an active unit in the TOE.

According to the Common Criteria, security attributes are tied statically to information, respectively to the objects it is linked to, whereas the security attributes of a subject may vary during runtime. Therefore, the control status of an object is invariable whilst the security level of a subject preliminarily receives an initial value.

The security attribute $C(O)$ receives the value “Strong” if it is located at a data depository that has to be controlled, otherwise, the security attribute $C(O)$ receives the value “Weak”. In this context a data depository D is called a **data depository requiring control** if the control flag CF is set to “True” in one of the most specific information flow rules for D .

Generating a new subject S (e.g. starting a new application), the security attribute $L(S)$ receives the value “Low”.

Example. Provided that the list of information flow rules $R1 - R3$ of Figure 1 are consistent we now consider an object O that is located at data depository $D3$. If $CF(R2) = \text{“True”}$ and $CF(R3) = \text{“False”}$, then it follows by definition that $C(O) = \text{“Strong”}$. This holds even in situations where the information flow rule $R3$ has to be applied.

Security Principles

The following security principles are defined for the discretionary information flow control policy.

- (P1) **Logging.** Decisions concerning the permission resp. denial of information flows are logged if this is requested by the information flow rules.
- (P2) **Data Security.** Authorised information flows always take place in accordance with the information flow instructions mentioned in the information flow rules.
- (P3) **Appropriation.** If the control status of an object O is “Strong”, information flows concerning object O are only allowed if they are requested by a subject S , which is authorised to do so according to the information flow rules.
- (P4) **Information Flow Control.** If the control status of an object O is “Strong”, an information I , deriving from object O may not be transferred to an object O^* with the control status “Weak” unless subject S triggering this information flow is authorised to do so according to the information flow rules.
- (P5) **User’s Discretion.** As an exception to principle (P4 – Information Flow Control), at least the TOE-administrator may authorise the information flow explicitly, i.e. an information I , deriving from object O with the control status “Strong” may then be transferred to an object O^* with the control status “Weak”.

The security principle (P4 – Information Flow Control) cannot be completely realized with the help of the security attribute control status, since the decision whether an information flow with operation ‘write’ is permitted depends on the origin of the information I . This explains the need for the security level assigned to subjects. The security attribute $L(S)$ is set to “High”, if subject S reads information I originating from an object O with control status “Strong”, except that subject S is considered to be trusted according to the actual information flow rule, i.e. the trust flag TF is “True”.

In addition to the TOE administrator other users may be allowed to carry out authorizations according to principle (P5 – User’s Discretion). Their permission may typically be based on certain owner rights.

Security Features

Let a list of control rules be specified. In order to decide whether a requested information flow is permitted or denied, it has to be checked first whether the data depository belonging to the requested information flow is mentioned in any information flow rule. If this is not the case, rules (R1) respectively (W1) are to be applied. Otherwise, before proceeding according to (R2) or (R3) respectively (W2) or (W3), it has to be checked first according to which of the existing information flow rules decisions have to be taken. A **selection function**, which is a crucial parameter of the discretionary information flow control policy, provides the actual information flow rule. If the protocol flag PF is “True” in the actual information flow rule the policy decision is being logged (cf. security principle (P1 – Logging)).

The discretionary information flow control policy is based on a number of technical rules summarized in Table 2. For the details of these rules we refer the reader to the upcoming protection profiles.

Selection function provides ...		no rule	rule R (P1)		
Control status		C(O) = Weak	C(O) = Weak	C(O) = Strong	
Subject S is ... (w.r.t. R)				appropriate	inappropriate
read		GO R1	GO (P2) R2	GO (P2&P3) R3 (i)	STOP (P3) R3 (ii)
write	L(S) = Low	GO (P4) W1 (i)	GO (P2&P4) W2 (i)	GO (P2&P3) W3 (i)	STOP (P3) W3 (ii)
	L(S) = High	STOP (P4&P5) W1 (ii)	STOP (P4&P5) W2 (ii)		

Table 2: Übersicht der Sicherheitsmaßnahmen (Security Features)

Limitations

There are some similarities to existing security policies like e.g. the well known Bell/LadPadula security model [3, chap. 7.3]. Roughly, one can think of control status as object classification and of security level as subject clearance. Moreover, it is possible to construct an access matrix out of our information flow rules. But such a matrix would typically fail to meet the “star property”, which is the central information flow principle in the Bell/LaPadula model. The reason is that the Bell/LaPadula matrix is a static construction expressing all possible read/write operations that subjects can perform on objects. In our discretionary information flow policy each decision is based on the dynamic behavior of the subjects under control, i.e. the control is based on all read/write operations that subjects actually perform on objects. In general, this approach is more flexible than the traditional security policies. On the other hand our policy currently supports only very elementary security attributes. The reason is that the focus of the presented work was not on the design of a sophisticated multi-level security policy, but on the formulation of the core principles and features of a discretionary information flow control policy. A considerably more powerful information flow policy can be constructed by incorporating hierarchical security attributes.

The appropriation principle of our security policy is of a very basic nature, since there is no requirement that controlled information stays within data depositories linked with a certain set of subjects. Again the reason is that we wanted to set up a core framework for the formulation of more sophisticated security principles. We currently ponder about the ways to extend the presented security policy in order to incorporate the information flow principle of the security policy of the British Medical Association [3, sec. 8.2.3]: “Information derived from record A may be appended to record B if and only if B’s access control list is contained in A’s.” Substituting the words ‘access control list’ with ‘set of subjects’ roughly demonstrates the correspondence with the presented work.

Conclusion

The presented discretionary information flow control policy is being incorporated in two related Common Criteria Protection Profiles which are scheduled to be certified and published by the German Information Security Agency (Bundesamt für Sicherheit in der Informationstechnik – BSI). These protection profiles provide the core principles and features of an extendable information flow control policy. They are influenced and driven by German federal and state data protection commissioners. It is planned to base certain technical parts of data protection regulations on new requirements including recommendations for the use products compliant with the presented protection profiles. Together with some promising international efforts we can hope for many products realizing privacy enhancing technologies (PETs) available in the near future.

Acknowledgements

We would to thank T-Systems ISS GmbH for evaluating the presented protection profiles and the following partners for their valuable comments on several drafts of the security policy:

- Bundesbeauftragter für den Datenschutz (German federal data protection commissioner)
- Der Bayerische Landesbeauftragte für den Datenschutz (data protection commissioner of the state of Bavaria)
- SuSE Linux AG
- Argus Systems Group [Germany] GmbH
- Utimaco Safeware AG
- Datapol GmbH

More information on the project and our partners as well as the protection profiles themselves are available at <http://dic.dfki.de/>.

Literature

- [1] Bundesamt für Sicherheit in der Informationstechnik (Ed.). *Discretionary Information Flow Control (SU)*. Common Criteria Protection Profile, to be published at <http://www.bsi.bund.de/>.
- [2] Bundesamt für Sicherheit in der Informationstechnik (Ed.). *Discretionary Information Flow Control (MU)*. Common Criteria Protection Profile, to be published at <http://www.bsi.bund.de/>.
- [3] R. Anderson. *Security Engineering*. Wiley Computer Publishing, 2001.