Protection Profiles for Remailer Mixes

Giovanni Iachello[1] and Kai Rannenberg[2]

[1] Telematics, IIG, Freiburg University
g.iachello@iol.it
[2] Microsoft Research Cambridge, UK
kair@microsoft.com

Abstract. In the past independent IT security evaluation according to published criteria has not realized its potential for the assessment of privacy enhancing technologies (PETs). The main reason for this was, that PETs were not covered appropriately in the evaluation criteria. This situation has changed somewhat, and therefore this paper reports on a case study, in which we developed Protection Profiles for remailer mixes. One reason for the development of these Protection Profiles was to test the privacy related components in the new Evaluation Criteria for IT Security – Common Criteria (International Standard 15408, ECITS/CC) and to develop improvements. Another reason was to contribute to an independent evaluation of privacy enhancing technologies. The experiment shows, that the ECITS/CC enable PPs for remailer mixes, but that there are still improvements necessary. The paper presents the Protection Profiles and the structured threat analysis for mixes, on which the Protection Profiles are based.

1 Introduction

Independent IT security evaluation can be very useful for privacy enhancing technologies (PETs), as PETs very often aim at the protection of individual users, and this is exactly the user group that usually does not have the resources to assess IT on its own. Of course evaluations and criteria then have to cover privacy aspects properly. This is not trivial, and early IT security evaluation criteria like the TCSEC and the ITSEC caught much criticism for their lack of coverage of privacy-related requirements, and for their tendency towards ever increasing data storage and centralization of trust. Meanwhile, evaluation criteria, like the recent Evaluation Criteria for IT Security – Common Criteria (International Standard 15408, ECITS/CC) contain components assigned to privacy. Therefore we used them to specify a number of Protection Profiles for remailer mixes. One reason for the development of these Protection Profiles was to test the privacy related components in the ECITS/CC and to develop improvements. Another reason was to contribute to an independent evaluation of privacy enhancing technologies.

The paper commences with an introduction into IT security certification and evaluation criteria (Chapter 2) and an overview of their problems regarding privacy and multilateral security (Chapter 3). It then describes the new ECITS/CC and their privacy components (Chapter 4). Chapters 5 and 6 describe the approach of writing PPs for remailer mixes and give a short introduction into mix technology. Chapter 7 presents the Protection Profiles and their rationales. Chapter 8 summarizes the experiences gained by writing the Protection Profiles. Chapter 9 proposes changes to the ECITS/CC. Chapter 10 gives a summary and conclusion. The Annex provides not only the references (Chapter 11) but also the three proposed functional families in a notation conformant with the prescriptions of the ECITS/CC (Chapter 12).

2 IT security certification and evaluation criteria

The complexity of today's information technology (IT) makes it impossible to evaluate its security by simple "examination". However, it is scarcely possible for many users to conduct more detailed checks, which are necessary for a qualified evaluation, as they cannot afford the expenditure this would entail. Thus, more and more users are faced with the problem of knowing very little about the technique they use for important transactions (e.g. processing sensitive patient data, signing documents, or making payments).

One way to enable confidence in IT is to evaluate and certify products and systems by neutral and competent institutions on the basis of published IT security evaluation criteria. Related certification schemes exist since the mid 80's, for example, in the USA, the UK, and Germany. There are differences between the schemes, but typically a sponsor asks (and pays) for an evaluation that is conducted by an accredited (commercial or governmental) IT Security Evaluation Facility (ITSEF) and monitored and certified by a (governmental or commercial) Certification Body (CB), cf. Figure 1. In most cases the sponsor of an evaluation is the manufacturer of the Target of Evaluation (TOE). An overview of Certification Schemes and more details can be found in [19].

To enable comparisons of evaluation results, criteria catalogs have been developed, which structure the IT security requirements. Some examples are given in Table 1.

While the TCSEC [22] had a rather fixed security model aiming at the confidentiality of military information, subsequent criteria e.g. the ITSEC [6] had a broader scope. These criteria are frameworks that IT manufacturers, vendors, or users can use to specify what security functions (Functionality) they wish to have evaluated and to what depth, scope, and rigor the evaluation should be performed (Assurance).

Functionality refers to the behavior of the product with regard to security concerns, while assurance allows stating requirements on e.g. the development process, the evaluation of the compliance to the requirements documents, the preservation of security during installation and maintenance, and the documentation. In practice, these requirements specify a series of actions, which the developer, the writers of documentation and the evaluators must complete.

Independent evaluation can be very useful for privacy enhancing technologies, as those very often aim at the protection of individual users, and this is exactly

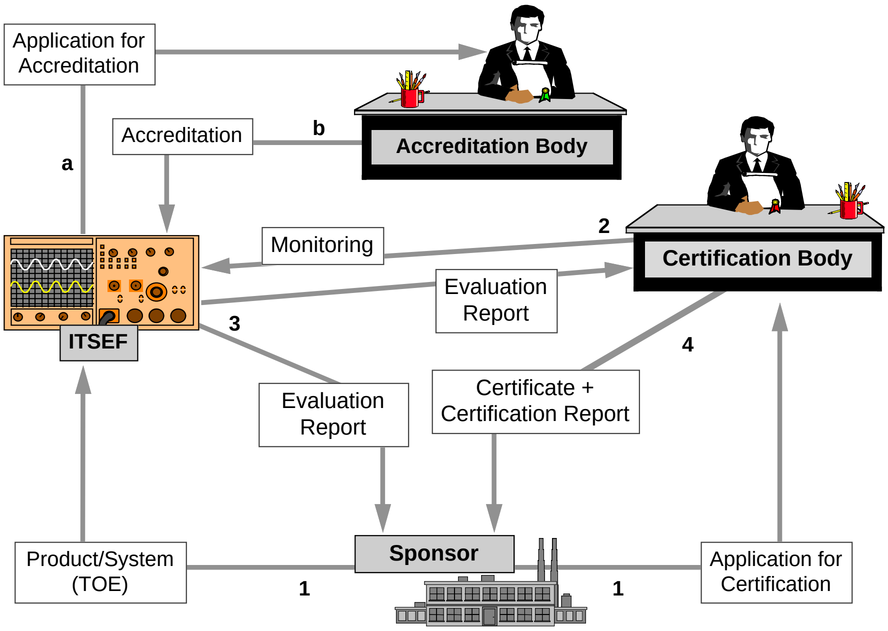

Fig. 1. Evaluation and certification of a TOE (1-4) and accreditation of an ITSEF (a-b)

Publication / Project Dates	**Editors**	**Criteria Name Current Version**
1983/85	USA Department of Defense (DoD)	Trusted Computer System Evaluation Criteria (TCSEC - "Orange Book")
1990/91	Commission of the European Communities (CEC)	Information Technology Security Evaluation Criteria (ITSEC) Version 1.2
1990-99	International Organization for Standardization / International Electrotechnical Commission ISO/IEC JTC1/SC27/WG3	Evaluation Criteria for IT Security (ECITS) International Standard 15408: 1999
1993-99	Common Criteria Project Govt. Agencies from CDN / D / F / GB / NL / USA	Common Criteria (CC) Version 2.1

Table 1. Some IT security evaluation criteria and their editors

the user group that usually does not have the resources to assess IT on its own. Of course evaluations and criteria then have to be comprehensive, especially regarding privacy. As the next chapter shows, this was not the case.

3 Problems regarding privacy and multilateral security

Various aspects of security certification and the underlying early criteria have been criticized, for example thebalance of the criteria and the meaningfulness and use of results (cf. e.g. [10, 18, 19]).

A main point of criticism from the application side was that the criteria were too biased towards hierarchically administered systems and the protection of system operators. The criteria seemed to not consider the fact that dangers are caused not only by users and outsiders, but also by operators and manufacturers of the systems. So there was a lack of *multilateral security*, i.e. taking the security requirements, not only of operators, but also of users and customers into account. Especially privacy aspects of telecommunication transactions were not covered, e.g. unobservability of calls to help lines or anonymous access to patent information on the Internet.

From a technical point of view systems with distributed organization and administration were only insufficiently covered. Also data-collecting functionality was overemphasized, while data economical functionality was ignored.

The following example illustrates how this lack of consideration for user protection in the criteria affects evaluation results. It also shows that the evaluation, which is described, was focused on the protection of the operators and neglected the protection of users or customers. A function for the selective logging of activities of individual users was classified as a non-critical mechanism that did not need evaluation. In the opinion of the evaluators, failure of this mechanism would not create weaknesses because if the function was not active, the activities of all users were logged [8]. From the operator point of view no real security risk existed, because no audit data would be lost - only perhaps more data than planned would be collected. However, from the users' point of view this is a considerable risk, because excessive logging and the resulting data can lead to substantial dangers for users and customers, e.g. when this data is misused.

4 The new ECITS/CC and their privacy components

Since 1990 two initiatives aim at globally uniform evaluation criteria, mainly to enable the mutual acknowledgement of evaluation results. A joint committee of ISO and IEC (JTC1/SC27/WG3) developed the "Evaluation Criteria for IT Security" (ECITS), which are being finished as IS 15408 [16]. In parallel to the ISO/IEC standardization, North American and European government agencies developed the "Common Criteria" (CC). Since CC Version 2.0 [3] there is a large convergence with the ISO ECITS, and CC Version 2.1 [4] and IS 15408 are fully aligned. After the problems with earlier criteria had also been brought up in ISO/IEC the new criteria contain a section aiming at privacy protection (cf.

Chapter 4.2). At the moment there are no plans for another version of the CC, but the ECITS will undergo the usual periodic revision of ISO/IEC standards, which will probably be done by JTC1/SC27/WG3 in 2003.

4.1 Overview of the ECITS/CC

The ECITS/CC share the goals and general approach of other evaluation criteria, as briefly introduced in Chapter 1, but provide a more flexible structure regarding functional and assurance requirements. In fact, they provide a catalogue of *functional requirements components*, which is a modular, structured library of customizable requirements, each of which tackles one specific aspect of the security requirements for the TOE. The Criteria provide also a catalogue of *assurance requirements*, which are grouped in seven ordered subsets, of increasing depth, scope and rigor.

On the one hand, these modifications create more liberty for the formulation of security targets, but on the other hand, they make the comparison of evaluation results more complicated. In order to resolve this problem and still give users the opportunity to formulate their own requirements, the CC introduced the concept of the "Protection Profile" (PP). A PP describes the functionality and assurance requirements for a certain application (e.g. health care administration) or technique (e.g. firewalls). Ideally, several products will be evaluated against a single PP, so that the results can be compared. ISO is setting up a regulated registry for PPs and the CC project is maintaining a PP list [5].

The ECITS/CC also provide a catalog of seven Evaluation Assurance Levels (EALs). These are an ordered set of packages of assurance components. Each EAL contains the lower level EAL and adds to it some other assurance requirements. The EALs are largely derived from the ITSEC. PP authors, who wish to concentrate on the functional requirements of their PP, can simply choose an EAL.

In most cases the ECITS/CC model the properties and behavior of a TOE by specifying a set of relevant entities and imposing constraints on the relationships between such entities. For this purpose, the following four entities are defined:

- Subject: "*An entity within the TSC*[1] *that causes operations to be performed*"; this can be, for example, a UNIX process;
- Object: "*An entity within the TSC that contains or receives information and upon which subjects perform operations*"; for example a file, a storage medium, a server system, a hardware component;
- Operation: a process initiated by a subject or user, which employs a subject to interact with one or more objects or subjects; this term is not directly defined in the criteria's' glossary.

[1] TSC = TSF Scope of Control: The complete set of interactions that are under the control of the TOE to satisfy its security requirements and to implement its security features.

– User: "*Any entity (human user or external IT entity) outside the TOE that interacts with the TOE*". It is necessary to clearly distinguish between "Subject" and "User". "User" is the physical user with all its attributes (name, role ...), or an external IT entity (i.e. another system interacting with the TOE), that initiates operations on the TOE, which are carried out, on its behalf, by "Subjects" operating in the TOE.

4.2 The ECITS/CC privacy families

The ECITS/CC contain four *Functional Families* directly related to privacy and organized in a privacy *Class*. Some of their components were inserted late in the criteria development process (for example, some of the Unobservability components were not present in version 1.0 of the CC). Most components have several levels, which sometimes are organized in hierarchies. A hierarchical level contains extra requirements. The following description of the components sticks close to that in the ECITS/CC.

Anonymity (FPR_ANO) Anonymity ensures that a user may use a resource or service without disclosing the user's identity. The requirements for Anonymity provide protection of the user identity, but Anonymity is not intended to protect the subject identity. There are two hierarchical levels:

FPR_ANO.1 Anonymity requires that other users or subjects are unable to determine the identity of a user bound to a subject or operation.
FPR_ANO.2 Anonymity without soliciting information enhances the requirements of FPR_ANO.1 by ensuring that the TSF does not ask for the user identity.

Applications include the ability to make inquiries of a confidential nature to public databases, respond to electronic polls, or make anonymous payments or donations.

Pseudonymity (FPR_PSE) Pseudonymity ensures that a user may use a resource or service without disclosing its user identity, but can still be accountable for that use. There are three partially hierarchical levels.

FPR_PSE.1 Pseudonymity requires that a set of users and/or subjects are unable to determine the identity of a user bound to a subject or operation, but that this user is still accountable for its actions.
FPR_PSE.2 Reversible pseudonymity requires the TSF to provide a capability to determine the original user identity based on a provided alias. FPR_PSE.2 is hierarchical to FPR_PSE.1.
FPR_PSE.3 Alias pseudonymity requires the TSF to follow certain construction rules for the alias to the user identity. FPR_PSE.3 is hierarchical to FPR_PSE.1.

Applications include the ability to charge callers for premium rate telephone services without disclosing their identity, or to be charged for the anonymous use of an electronic payment system.

Unlinkability (FPR_UNL) Unlinkability ensures that a user may make multiple uses of resources or services without others being able to link these uses together.

> *FPR_UNL.1 Unlinkability* requires that users and/or subjects are unable to determine whether the same user caused certain specific operations in the system.

Applications include the ability to make multiple use of a pseudonym without creating a usage pattern that might disclose the user's identity.

Unobservability (FPR_UNO) Unobservability ensures that a user may use a resource or service without others, especially third parties, being able to observe that the resource or service is being used. There are four partially hierarchical levels.

> *FPR_UNO.1 Unobservability* requires that users and/or subjects cannot determine whether an operation is being performed.
> *FPR_UNO.2 Allocation of information impacting unobservability* requires that the TSF provide specific mechanisms to avoid the concentration of privacy related information within the TOE. Such concentrations might impact unobservability if a security compromise occurs. FPR_UNO.2 is hierarchical to FPR_UNO.1.
> *FPR_UNO.3 Unobservability without soliciting information* requires that the TSF does not try to obtain privacy related information that might be used to compromise unobservability.
> *FPR_UNO.4 Authorised user observability* requires the TSF to provide one or more authorized users with a capability to observe the usage of resources and/or services.

Applications include technology for telecommunications privacy, especially for avoiding traffic analysis to enforce constitutional rights, organizational policies, or defense requirements.

5 Experimenting by writing Protection Profiles for mixes

As the ECITS/CC aim at covering security requirements also for untraditional applications that were not covered in earlier criteria, it seemed useful to experiment with the new criteria by using it. Actually, during the development of the CC a number of example PPs had been produced on the basis of CC V1.0 for testing purposes (see [5]), but there was no (published) PP aiming at privacy requirements. To gain more experience as to whether the ECITS/CC are complete and adequate enough to express requirements on privacy friendly functionality, some example PPs were written.

The mix application (cf. Chapter 6) was chosen because it is a prime example of a distributed application where multilateral security concerns involving operators and users come into existence. The availability of an extensive literature on the subject, of real world implementations and the interest that anonymous and untraceable communication have gained recently, are also all favorable reasons which make this kind of application an ideal testing ground.

The development of the PPs started with an initial survey of the mix literature and implementations to get acquainted with the underlying concepts and technology. A mix implementation was installed and operated in a controlled manner for a couple of weeks. This process resulted in the enumeration of a set of threats that were then used as the basis of the PPs (cf. 7.2).

6 Short introduction into mixes

A mix is a remailer system with the objective of hiding the correspondence between sender and recipient of a message. The concept was introduced by D. Chaum in 1981 [7] and has been subsequently refined and applied to other applications besides email, such as ISDN telephony and WWW access (e.g. [21, 17, 24]). This basic functionality allows achieving unlinkability of communicating partners, but anonymity can also be achieved if the sender does not explicitly state its identity in the message. As a further development also pseudonymity can be implemented using a mix remailer system, using so-called "return addresses".

There are at least two working implementations of mixes: the first one, is a free software called Mixmaster [9], which evolved from a first-generation plain anonymizing remailer to a complete mix system in 1994. The software is now in use at various sites, but is generally administered by volunteers and thus not apt for widespread commercial use.

A commercial pseudonym-based mix system is being produced by Zero Knowledge Systems [24], which offers a client product for sending email through a set of independently administered nodes. Some of these nodes are administered by ZKS, which also produces the remailer software.

A mix system achieves untraceability of messages essentially by deploying a distributed architecture, where each node is independently administered. The sender selects a path in the mix network to reach the receiver, and each node resends the message to the next one according to instructions present in the message itself. The message is encrypted in such a way that each relay node only gets to know the node from which it received the message and the node to which it forwarded the message.

7 The Protection Profiles written

Several Protection Profiles were written to cover the features and threats regarding mixes and to test the Common Criteria privacy components. Section 7.1 documents the development history of the PPs and their versions, Section 7.2 gives an overview of the threats considered. The remaining sections in this chapter document the three PPs:

- "Single Mix PP" (7.3);
- "Protection Profile for an Unobservable Message Delivery Application Using Mixes" (or "Multiple Mix PP") (7.4);
- "User-oriented PP for Unobservable Message Delivery Using Mix Networks" (7.5) This PP is described in most detail giving a mapping from threats and other assumptions to security objectives and functional components.

7.1 Development history of the Protection Profiles

The development history of the Protection Profiles already shows some of the issues coming up and is therefore documented here. The following figure gives an overview of the development history; the boxes represent PPs and are positioned in a temporal order (temporal axis from top to bottom); the arrows connecting the PPs represent a "flow of knowledge" (i.e. e. addressed threats, security objectives) from one PP to the other.

Initially, a choice was made to write two PPs following an "architectural" subdivision suggested also by the criteria, i.e. writing a PP for a single mix node ("Single Mix Protection Profile"), and then one for the whole mix network ("Protection Profile for an Unobservable Message Delivery Application Using Mixes", or "Multiple Mix PP" [12]). This represents a rather traditional way of subdividing security problems into manageable pieces, and derives from the "secure each part to secure the system" paradigm.

In the case of the mix network, however, this path resulted in a dead-end, because the second PP, which stated requirements for the whole network, actually tried to make a compromise between the security requirements of the network, and those of the user of the network. For example, a standard option for protecting the mix network from some kinds of flooding attacks is that of keeping audit logs. This clearly endangers the potentially anonymous users, because corrupt administrators or others gaining access to the logs could use the information contained therein to trace back the messages to their original senders.

After unsuccessfully following this path we decided for an alternative approach. This was to divide the security requirements documents based on the so-called "multiple interests" paradigm, i.e. writing one document for each of the involved parties, which in the mix system are the administrators and the users, and each time taking into account the security needs and concerns of the focused party. This approach again led to the writing of two documents: the "Single Mix Protection Profile" [11], which was largely rewritten from the first PP with the same title, and the "User-Oriented Protection Profile for unobservable message

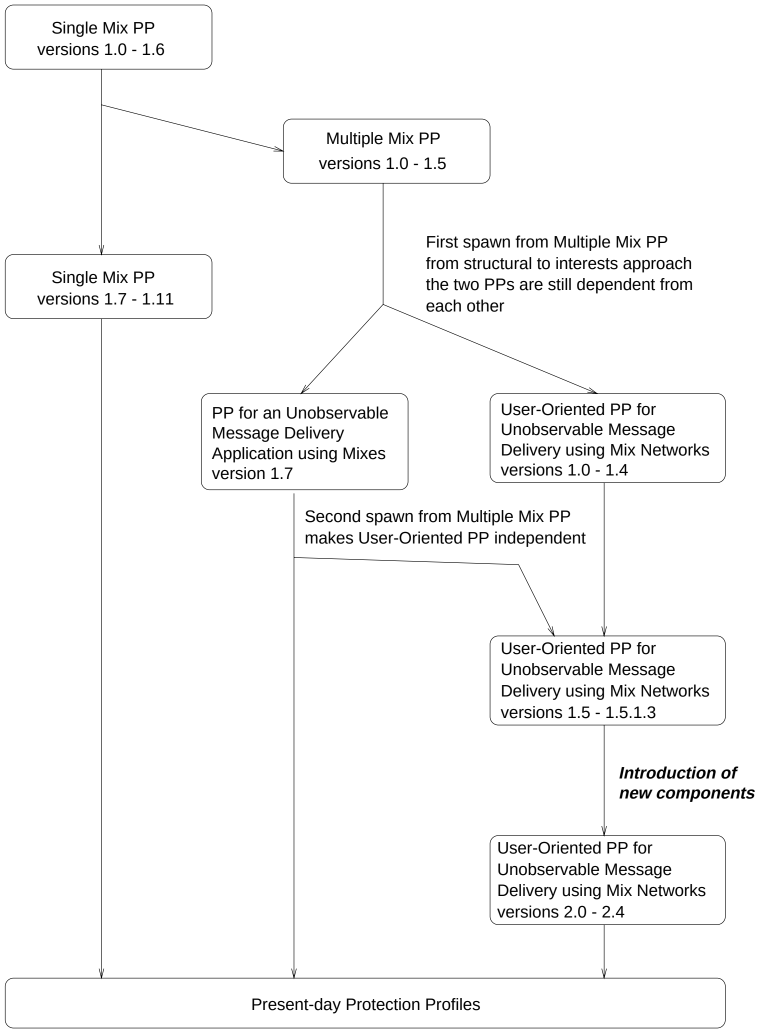

Fig. 2. PP Development history

delivery using Mix networks" [13], which incorporated parts of the "Multiple Mix PP".

The former document states the security requirements of a single node, which overlap largely with the requirements as felt by the administrator of such node (e.g. resistance to attacks, secure operating environment, physical protection ...), while the latter addresses the needs of the user with respect to the whole network, and includes requirements of anonymity, unlinkability of communicating parties, etc.

Eventually, it was found that the main challenges for the expressive power of the Criteria were posed by this second document, because some of the security requirements related to fundamental privacy-enhancing properties were not to be found in the stock ECITS/CC components (cf. Chapter 8).

Choosing an EAL (see section 4.1) was easier than formulating the functional requirements. The choice is influenced by many external factors, which include the intended use and operational environment of the TOE, policies of the organization that will deploy the TOE, and the will of the sponsor to let the product be evaluated at a high level (which rises the evaluation costs).

Two rather different choices were made: for the "Single Mix PP" and the "Multiple Mix PP" a relatively low level (EAL 3) of assurance was requested; this choice was justified by the fact that the mix is a rather simple system, where the architectural security strengths derive not from the single system, but from the fact that multiple systems operate together.

The "User-Oriented Mix PP" follows another approach. The idea in this case is that the user wants to gain full assurance that the single mix systems were correctly developed, and that the architecture and project, as a whole, were examined by independent organizations. EAL 5 was chosen because it is the first EAL that introduces *complete independent testing* of the TOE.

It is however to be noted that an independent test of the TOE is not sufficient to assure the user that the system will not be malevolently administered after deployment. The ECITS/CC assurance requirements did not aim at evaluating the operation of deployed systems. Closer to this task are risk management standards like IS 13335 [15] or BS7799 [1] and related certification schemes like c:cure [2].

7.2 The threats considered in the Protection Profiles

Considering and documenting threats to a TOE is the basis of a PP. The threat list must be *complete* and *relevant*. Obviously, there is no guarantee, that a list of threats is complete. Therefore peer review and multiple incremental cycles are necessary. Each PP was rewritten many times before reaching a stable state for the time being. Additionally, the threats must be stated in a manner to ease the formal demonstration of correspondence with the security objectives, to the degree mandated by the choice of the EAL.

The threat lists are summarized in Table 2, where they are subdivided according to the three Protection Profiles written for the mix system. The threats are briefly described in terms of implementation, effects and countermeasures, and ordered by type. The threat type isone of:

- Active: the threat requires an attacker to actively interfere with the operation of the mix or network, e.g. by blocking communications,
- Passive: this kind of threat is limited to passive operation (e.g. observing traffic through a mix),
- Error: these threats derive from erroneous operation of the mix, due to e.g. bad configuration, etc.

The threats in the preceding list are then stated in each Protection Profile as formal threats, adhering to the requirements imposed by the PP structure as described in the CC, as shown in the following sections.

The complete text of the PPs [11–13] is freely available and also contains extensive justifications for the selection of threats and countermeasures.

7.3 Single Mix Protection Profile

The Single Mix Protection Profile was written to address the security problems of a single mix system, without consideration towards the necessities of the user (who wants to send anonymous mail) and also ignoring all the security threats, which may derive from the connection of the system with other mixes. The threat list of this Protection Profile includes items such as *flooding attacks, logical access to the TOE, replay attacks, traffic size analysis*, as shown in Table 3. The last two threats (marked with "TE") are intended to be countered not only by the mix itself but also by the environment (operating system, etc.).

At this point a general observation regarding the following lists of threats is necessary. These threat lists derive from a detailed analysis of the operation, and risks, of mix networks, both from a practical point of view and from a theoretical one. Afterwards, an informal threat list is produced (see the previous section), which is then used to build a more structured threat list, which complies with the structural requirements (ease of correspondence demonstration, avoiding overlapping threats, etc.) needed by the PP.

The lists must be, obviously, considered together with assumptions, which are also included in separate tables. The idea behind this structure is that the PP should aim at *completeness* in addressing the security issues, either by stating assumptions, or by indicating possible threats. However, the decision of how to subdivide assumptions and threats is a very delicate one, because assumptions clearly do not need to be addressed by the PP, but may hide some major security issues, thus causing the PP to be ineffective.

Table 4 lists the assumptions related to the previous threat list.

Table 5 shows the list of functional components used by the PP to address the shown threats. All the functional components are taken from the ECITS/CC catalog. Each of the selected components listed in the table introduces into the PP a number of atomic requirements that can be tailored by the author.

Name	**Type**	**Implementation**	**Permits Analysis ...**	**(Potential) Effects**	**Counter-measure(s)**
Single mix threats (threats to a single mix system, as seen by the mix operator, and basis for the Single mix PP)					
Logical access	Active	Gain access to the TSF data and algorithms	Of administrative logs Of mix operation	Total failure of mix security functions	Trusted OS, Limit remote administration
Physical access	Active	Gain access to the TSF physical location	Of administrative logs Of mix operation	Total failure of mix security functions	Trusted site
Administrator corruption	Active	Corrupt the administrator	Of administrative logs Of mix operation	Total failure of mix security functions	Organizational policies, Operation review
Replay attack	Active	Intercept and resend a message many times	Of outgoing traffic	Leak of (partial) information	Replay detection on message paths
Flooding attack (DoS)	Active	Flood mix with dummy messages	n/a	Interruption of service	Flooding resistant mix and OS, Origin check
Flooding attack (traffic analysis)	Active	Flood mix with known messages	Of outgoing traffic	Leak of information on singled-out	Origin check message path
Size analysis	Passive	Intercept messages and record their sizes	Of ingoing and outgoing traffic	Leak of information on message paths	Standard and fixed message size
Timing analysis	Passive	Intercept messages and record their transmission times	Of ingoing and outgoing traffic	Leak of (partial) information on message paths	Random delay strategies
Order analysis	Passive	Intercept messages and record their order	Of ingoing and outgoing traffic	(Partial) information on message paths	Random reordering strategies
Content-based traffic analysis	Passive	Intercept messages and read their content	Of ingoing and outgoing traffic	Leak of complete information on message paths	Encryption of message traffic
Mismanagement	Error	Mismanagement of some TSF	n/a	Loss of TOE security properties	Documentation, Design for manageability, Organizational policies
Processing error	Error	Accidental processing error resulting in truncation, loss, alteration of messages	n/a	Unreliable service	Redundancy assurance techniques

Table 2. Threats used as basis for the Protection Profile (part 1)

Name	Type	Implementation	Permits Analysis ...	(Potential) Effects	Counter-measure(s)
Multiple mix threats (threats to the entire mix network, or related to the network connections, and basis for the Multiple mix PP)					
Network block	Active	Block the network connections to part of the TOE	n/a	Interruption of service, Degraded service	Organizational policies, distribution of the TOE
Impersonation	Active	Intercept and redirect the network connections to part of the TOE	Of requested traffic in the impersonated network	Degraded service, Leak of information on message paths	Encryption, Sound key distribution policy
Message interception	Passive	Intercept and read messages	Message content exchanged by parts of the TOE	Leak of information on message paths	Encryption
Network unreliability	Error	Accidental damage of messages (truncation, loss, alteration)	n/a	Unreliable service	Redundancy (multiple path ...), Error detection and report
Mismanagement of network security functions	Error	Erroneous configuration of the TSF	n/a	Loss of security properties	Documentation, Design for manageability, Organizational practices
User mix threats (Threats as seen by the User and basis of the User-oriented mix PP)					
Untrusted mix	Active	A mix in the network may be compromised and reveal tracing information	Of transiting messages	Exposure of linking information	Division of trust
Mix conspiracy	Active	Some mixes in the network may conspire to share and analyze traffic information	Of transfer logs and TSF operation	Loss of expected security functionalities	Organizational policies, Independent administration
Forgery	Active	An attacker may send forged messages using a user's origin credentials	n/a	Loss of accountability properties	Use of digital signatures
Intercept	Passive	Messages are intercepted while transiting from user to a mix	Of incoming and outgoing traffic	Information on use patterns	Generation of dummy messages by the users
Misuse	Error	Erroneous use of the TSF by the user	n/a	Loss of expected security functionalities	Documentation, Ease of use
Unreliability	Error	The connecting network may be unreliable, resulting in message loss, truncation or alteration	n/a	Unreliable service	Redundancy, Error detection

Table 2. Threats used as basis for the Protection Profile (part 2)

Threat Label	**Description**
T.DenialOfService	An attacker may try flooding the mix with a great amount of messages, thus causing an overload on the mix and possibly leading to Denial of Service by the mix.
T.FloodingAttack	An attacker may try flooding the mix with a great amount of messages, to single out only one unknown message and discover its destination.
T.ForgeOrigin	An attacker may send to the mix messages with a forged origin field. *This can be done for various reasons, for example to hide a flooding attack.*
T.InterceptMessages	An attacker may intercept and read the content of the messages (including the message origin and final destination, arrival and departure order and time) exchanged by users and the mix or between mixes. *The attacker may use intercepted messages to perform a traffic analysis to reveal input/output message flow patterns.*
T.LogicalAccess	An attacker (this includes unauthorized users of the host system) may gain access to the mix. *This may cause the complete failure of the mix.*
T.ReplayAttack	An attacker may intercept an incoming message and feed it to the mix many times, and, by checking the outgoing messages, discover where it is directed.
T.SizeAnalysis	An attacker may track the sizes of incoming and outgoing messages, thus linking origin and destination.
T.WrapAndImpede	An attacker may completely wrap the mix, thus selectively or totally impeding exchange of messages with the users or other mixes.
TE.Untrustworthy Administrator	The mix administrator may abuse his trust and compromise completely the operation of the mix. *Possible actions include: recompiling the mix application and modifying its behaviour, so to trace messages, and impairing the mix and causing DoS.*
TE.Proper Administration	The mix may be administered by the mix administrator in an insecure or careless manner. *This includes both the administration of the mix itself, such as unintentionally disclosing the confidential security attributes of the mix, and the administrative practice, such as not using a trusted channel when remotely administering the mix.*

Table 3. Threats in the Single Mix Protection Profile

Assumption Label	Description
A.Environment	The mix works in a networked environment, on a single host.
A.Spam	In case of a spam attack, the mix may not be able to satisfy ist goal.
A.DedicatedHost	The mix is the only process on its host system. Its administrator coincides with the host's system administrator. *This assumption, and the following, is one possible formulation of the main relevant assumption: that the host operating system will not allow or cause security breaches against the TOE. Having the mix as the only application of the host system greatly reduces the complexity of the host's analysis.*
A.OS	The Operating System of the host of the mix identifies authorized users and protects the mix operation and its data with regard to confidentiality, integrity, availability, and accountability.

Table 4. Assumptions in the Single Mix Protection Profile

The selected EAL (Evaluation Assurance Level) is 3. This EAL was selected because it is commonly considered the highest attainable EAL through current not security oriented development practices. Moreover EAL 3 was considered as a good compromise between the TOE analysis complexity and the intended use of the TOE. Recall that a single mix is to be used in a network, and the real strength of the system relies upon the existence of a large number of independent systems.

7.4 Multiple Mix Protection Profile

The "Protection Profile for an Unobservable Message Delivery Application Using Mixes" (or Multiple Mix Protection Profile) was written initially to complement the previous, and to take into account both the entire network of mixes, and the requirements set by the user, which sees the mix network as one homogeneous and opaque entity. Thus, the threats this PP addresses include threats like *message interception* and *denial of service*, as shown in Table 6.

Some of the threats may appear to be too obvious to be included in the threat list (as the T.MixPeek threat, which states the possibility of a mix to read the information contained in a message which is not encrypted.) However, such a statement is necessary exactly to make sure that all messages which transit through the mix system are encrypted in such a way that each mix will not be able to read the content of the message apart from the information of the next node where to send it.

The list of related assumptions follows (Table 7). Some of the assumptions are stated only to simplify the PPs, like the A.DedicatedHost, which excludes other processes on the same host of each mix, and are really not essential. However, there are assumptions, like the A.MinimalTrust, which are very important, because they state explicitly when the entire mix network fails.

Short name	Unique name	Component Description and Comments
FAU_ARP.1	Security alarms	This family defines the response to be taken in case of detected events indicative of a potential security violation. *This requirement states what the mix should do when a security violation is detected (e.g. Spam attack, Access to the security functions . . .) For example, the TOE may inform the administrator of potential attacks, and possibly switch to a secure fail mode upon detection of security violations.*
FAU_GEN.1	Audit data generation	Audit data generation defines the level of auditable events, and specifies the list of data that shall be recorded in each record. *This component requires the TOE to generate an audit trail, which can then be used by an automated or manual attack analysis and by an attack alarm system.*
FAU_SAA.3	Simple attack heuristics	Simple attack heuristics, the TSF shall be able to detect the occurrence of signature events that represent a significant threat to TSP enforcement. This search for signature events may occur in real-time or during a post-collection batch-mode analysis. *This component is used to require the TOE to provide some means for automatic detection of (and thus reaction to) potential attacks.*
FAU_SAR.1	Audit review	Audit review provides the capability to read information from the audit records. *This component ensures that the audit trail is readable and understandable.*
FAU_STG.1	Protected audit data trail storage	Protected audit trail storage, requirements are placed on the audit trail. It will be protected from unauthorized deletion and/or modification. *This component is chosen to ensure that the audit trail is protected from tampering. Only the authorized administrator is permitted to do anything to the audit trail.*
FAU_STG.4	Prevention of audit data loss	Prevention of audit data loss specifies actions in case the audit trail is full. *This component ensures that the authorized administrator will be informed and will be able to take care of the audit trail should it become full. But this component also ensures that no other auditable events as defined in FAU_GEN.1 occur. Thus the authorized administrator is permitted to perform potentially auditable actions though these events will not be recorded until the audit trail is restore to a non-full status.*
FCS_CKM.1	Cryptographic key generation	Cryptographic key generation requires cryptographic keys to be generated in accordance with a specified algorithm and key sizes that can be based on an assigned standard. *This and the following two requirements are included in the PP but left unspecified, since cryptographic standards evolve rapidly.*
FCS_CKM.2	Cryptographic key distribution	Cryptographic key distribution requires cryptographic keys to be distributed in accordance with a specified distribution method that can be based on an assigned standard.
FCS_CKM.4	Cryptographic key destruction	Cryptographic key destruction requires cryptographic keys to be destroyed in accordance with a specified destruction method that can be based on an assigned standard.
FCS_COP.1	Cryptographic operation	Cryptographic operation requires a cryptographic operation to be performed in accordance with a specified algorithm and with a cryptographic key of specified sizes. The specified algorithm and cryptographic key sizes can be based on an assigned standard. *The type and strength of the cryptographic functions is also left unspecified, and must be determined in accordance to the intended use of the TOE, and the perceived threats.*
FDP_IFC.1	Subset information flow control	Subset information flow control requires that each identified information flow control SFP be in place for a subset of the possible operations on a subset of information flows in the TOE. *This requirement (and the following) identifies the security attributes (e.g. routing information) and the allowed information flows through the mix.*

Table 5. Functional Components in the Single Mix Protection Profile (part 1)

Short name	**Unique name**	**Component Description and Comments**
FDP_IFF.1	Simple security attributes	This component requires security attributes on information and on subjects that cause that information to flow or that act as recipients of that information. It specifies the rules that must be enforced by the function, and describes how security attributes are derived by the function.
FDP_IFF.4	Partial elimination of illicit information flows	Partial elimination of illicit information flows requires the SFP to cover the elimination of some (but not necessarily all) illicit information flows. *Information about the correlation of origin with destination may reach the attacker through a covert timing or storage covert channel, if care is not used in to blocking such information leakage; this information may help timing, order, and size analysis attacks, and flooding attacks. This requirement ensures that such leakage does not take place.*
FDP_RIP.2	Full residual information protection	Full residual information protection requires that the TSF ensure that any residual information content of any resources is unavailable to all objects upon the resource's allocation or deallocation. *This component requires the TOE not to retain data that could be used by an unauthorized user of the security attributes management functions or by a malevolent administrator to trace messages. (According to the ECITS/CC, this component only relates to "residual" data - storage space that is not overwritten after use, etc.)*
FDP_SDI.2	Stored data integrity monitoring and action	Stored data integrity monitoring and action adds the additional capability to the first component by allowing for actions to be taken as a result of an error detection. *This component is needed for the correct operation of the TOE. If a message is modified while out of TSF control (e.g. by an attacker), this component shall ensure that the message will be discarded as invalid prior to processing.*
FDP_UCT.1	Inter-TSF user data confidentiality transfer protection	Basic data exchange confidentiality, the goal is to provide protection from disclosure of user data while in transit. *This component ensures the data confidentiality during transport of the user data (namely, messages) between separate TSFs and between user and TSF. The FDP_UCT.1 and the FCS_COP.1 components work together (that is, the former requires messages to be confidential (e.g. by using encryption), the latter sets requirements on the cryptographic functions.)*
FMT_MSA.1	Management of security attributes	Management of security attributes allows authorized users (roles) to manage the specified security attributes. *Nobody may modify or change the security attributes associated with messages, as they are integral part of the data needed by the mix to operate correctly. The mix does not store user data other than the transiting messages, so there is no further data to manage.*
FMT_MSA.2	Secure security attributes	Secure security attributes ensures that values assigned to security attributes are valid with respect to the secure state. *This component requires the TOE to perform validity checks on the security attributes used by the TOE itself, such as (local) origin and destination addresses of messages, message signatures and keys, and the like.*
FPR_ANO.2	Anonymity without soliciting information	Anonymity without soliciting information enhances the requirements of FPR_ANO.1 by ensuring that the TSF does not ask for the user identity. *This component (and the following) ensures that the TOE can be used without the user being required of disclosing his own identity.*

Table 5. Functional Components in the Single Mix Protection Profile (part 2)

Short name	Unique name	Component Description and Comments
FPR_UNL.1	Unlinkability	Unlinkability requires that users and/or subjects are unable to determine whether the same user caused certain specific operations in the system.
FPT_FLS.1	Fail with preservation of secure state	Failure with preservation of secure state requires that the TSF preserve a secure state in the face of the identified failures. *This component is used to force the TOE into biasing its operations towards a more secure than reliable operation. The rationale behind this is that a user is more interested in using a safe mix, rather than reliable one, since the TOE is anyhow intended to be used in an environment where multiple mixes are in operation.*
FPT_RCV.1	Manual recovery	Manual recovery, allows a TOE to only provide mechanisms that involve human intervention to return to a secure state. *This component allows for administrators to restore mix operation after a failure; prior to reactivating the mix, however, the administrator shall analyze the audit records, understand the reason that caused the failure, and remove its cause.*
FPT_RPL.1	Replay detection	Replay detection requires that the TSF shall be able to detect the replay of identified entities.
FPT_STM.1	Reliable time stamps	Reliable time stamps requires that the TSF provide reliable time stamps for TSF functions. *This component (and the following) is selected to satisfy dependencies by other components.*
FPT_TST.1	TSF testing	TSF testing provides the ability to test the TSF's correct operation. These tests may be performed at start-up, periodically, at the request of the authorized user, or when other conditions are met. It also provides the ability to verify the integrity of TSF data and executable code.
FTP_ITC.1	Inter-TSF trusted channel	Inter-TSF trusted channel requires that the TSF provide a trusted communication channel between itself and another trusted IT product. *This component is selected to ensure the presence of a trusted channel in inter-TSF communication. The channel provides for confidential and untampered communication between trusted IT products, namely, mixes; such channel might not be reliable, nor does it provide for party identification.*
FTP_TRP.1	Trusted path	Trusted path requires that a trusted path between the TSF and a user be provided for a set of events defined by a PP/ST author. The user and/or the TSF may have the ability to initiate the trusted path. *This component is selected to ensure the presence of a trusted path between the TSF and the user; such a path might not be reliable, nor does it provide for identification of the communicating party.*

Table 5. Functional Components in the Single Mix Protection Profile (part 3)

Threat Label	**Description**
T.DenialOfService	The TOE may be isolated from the users by blocking the network connections, and causing DoS. *This threat applies to the TOE as well as to the surrounding environment. The PP will however only address if from the TOE point of view.*
T.MessageInterception	The network and physical layer connections between mixes are not trusted. *This means that an attacker may manage to intercept messages transiting over the network and read their origin and destination fields.*
T.Misuse	Users may improperly use the TOE and produce traceable messages, while thinking the message was correctly sent and delivered. The administrators may inadvertently mismanage or badly configure parts of the TOE as to loose the security properties of that part of the TOE.
T.MixPeek	A subverted mix may be able to gain knowledge of the origin and destination of a message by reading its content while processing it.
T.OneStepPath	A mix may gain information linking origin and destination if the path from the origin user to the destination user contains only one mix.
T.TOESubstitution	An attacker may block messages sent by some user and act as the TOE, or a part thereof. Inadvertent users may send messages to the attacker instead of to the TOE, and the attacker may then read origin and destination data and forward the message to the destination.
T.UnreliableNetwork	The connecting network may not be reliable on correctly delivering messages between parts of the TOE. Specifically, messages may be lost, altered or truncated accidentally.
TE.MixConspiracy	Subverted mixes may share input/output information with the goal of linking origin and destination of a message.

Table 6. Threats in the Multiple Mix Protection Profile

Assumption Label	**Description**
A.IndependentAdministration	The mixes forming the TOE are assumed to be independently administered from each other.
A.MinimalTrust	The TOE may not be able to reach its goal if all nodes (mixes) are subverted.
A.OpenEnvironment	The mix network works in an open networked environment; each mix is operated on a single host.
A.UserCooperation	Users cooperate actively at the enforcement of the security policy of the TOE. *Users are trusted to use in a correct manner the services made available by the TOE to reach their anonymity goals.*
A.DedicatedHost	The mix is the only process on its host system. Its administrator coincides with the host's system administrator.
A.SecureLocation	The mixes forming the TOE are located at secure sites and physically protected from access by unauthorized users.

Table 7. Assumptions in the Multiple Mix Protection Profile

This document tries to conciliate the needs of the operators of the mixes on the network with those of the users, and this leads to a conflict, which is difficult to solve using the standard CC components. Table 8 shows the components used to specify the requirements for this PP.

Unless otherwise indicated, the components are described and commented on similarly to the corresponding components of the Single Mix PP (cf. Table 5).

The selected EAL is 3, because the higher EALs are mainly focused on the enhancement of the development process, while in the case of the PP the development is of secondary importance with respect to the installation and operation of the system.

7.5 User-oriented Protection Profile for Unobservable Message Delivery Using Mix networks

The "User-Oriented Protection Profile for Unobservable Message Delivery Using Mix Networks" was developed with in mind only the needs of the user of the mix network, and thus addresses threats like *untrusted mix, misuse, key forgery*, as shown in Table 9. The table also includes two Organisational Policies (marked with an "O." label) that state supplementary requirements for the TOE and that do not derive directly from some threat. The policies are however treated like threats in the following steps that lead to the formal requirements statement.

A set of assumptions for the User-Oriented PP follows in Table 10.

The second step of writing a PP is that of specifying a set of security objectives, which state the objectives that the TOE should reach to be able to counter all the threats. Table 11 shows the Security Objectives that were stated for this PP. As for the threats table, the Security Objectives are also divided in two categories, namely, objectives which are to be achieved solely by the TOE, and objectives for which the surrounding environment (Operating System, Administration, etc.) are partly or wholly responsible.

A correspondence between objectives and threats must be demonstrated (the rigor of this analysis depends on the EAL selected for the PP), but such demonstration is omitted here due to space constraints. However, the correspondence between threats and objectives is shown in Table 12.

As written above (section 7.2), the problems encountered during the development of this PP because of expressive deficiencies of the CC components led eventually to the writing of the proposed families. After the development of the new components, a second version of the PP was written; this new PP maintains the same threats and objectives of the previous PP, and simply uses also the new components to express its requirements, accordingly to the recommended top-down practice for PP development. The new version of this PP is decidedly simpler, more effective, and more precise in the requirements definition for the considered application. Table 13 shows the functional components used by this new version of the PP, which also employs the new proposed components (marked in the third column). Where relevant, a short description of the component and its use is provided in the fourth column.

Short name	**Unique name**	**Component Description and comments**
FCS_CKM.1	Cryptographic key generation	
FCS_CKM.2	Cryptographic key distribution	
FCS_CKM.4	Cryptographic key destruction	
FCS_COP.1	Cryptographic Operation	
FDP_IFC.1	Subset information flow control	This component requires that each identified information flow control SFP be in place for a subset of the possible operations on a subset of information flows in the TOE. *This component defines the policy of operation of the TOE and the subjects, information and operations controlled by the TOE.*
FDP_ITT.1	Basic internal transfer protection	Basic internal transfer protection requires that user data be protected when transmitted between parts of the TOE.
FDP_ITT.3	Integrity monitoring	Integrity monitoring requires that the SF monitor user data transmitted between parts of the TOE for identified integrity errors. *This component is required to allow safe delivery of messages through the mix network.*
FDP_RIP.2	Full residual information protection	
FMT_MSA.1	Management of security attributes	
FMT_MSA.2	Secure security attributes	
FMT_MSA.3	Static attribute initialisation	Static attribute initialisation ensures that the default values of security attributes are appropriately either permissive or restrictive in nature. *The security attributes (hash values, signatures ...) of the data stored and transferred throughout the TSF are generated automatically by the TOE. This data is not discretionary in nature, but must obey specific rules and may not be changed by users, or by the mix administrator.*
FPR_ANO.2	Anonymity without soliciting information	

Table 8. Functional Components in the Multiple Mix Protection Profile (part 1)

Short name	**Unique name**	**Component Description and comments**
FPR_UNL.1 (1)	Unlinkability of origin and destination	Unlinkability requires that users and/or subjects are unable to determine whether the same user caused certain specific operations in the system. *This component is introduced here to make sure that the network actually will grant the unlinkability of origin and destination of a message.*
FPR_UNL.1 (2)	Unlinkability/ untraceability	*This requirement is stated to make sure that an observer may not be able to link two observed messages transiting through the mix network, as being steps of the same message chain. This somewhat awkward formulation of the unlinkability requirements simply states that a mix shall not be able to bind messages exchanges between other nodes together into a single mix chain.*
FPR_UNO.2	Allocation of information impacting unobservability	Allocation of information impacting unobservability requires that the TSF provide specific mechanisms to avoid the concentration of privacy related Information within the TOE. Such concentrations might impact unobservability if a security compromise occurs. *Particularly, this requirement states that routing information may be accessible to mixes only when strictly necessary, e.g. to identify the following step in the mix chain as described for example in [7]. This functional component provides protection both to the mix network, by minimizing the exposure of information to attackers, which may be used to exploit covert channels, and to the user, to guarantee that the network will continue to operate securely even when some, unless not all, nodes are compromised.*
FPT_FLS.1	Failure with preservation of secure state	*If some nodes in the network fail or are subverted, the remaining nodes shall continue to work properly, in a secure manner.*
FPT_ITT.1	Basic internal TSF data transfer protection	Basic internal TSF data transfer protection, requires that TSF data be protected when transmitted between separate parts of the TOE. *This component (and the following) protect the data produced and used by the TSF, and transferred between parts of the TOE, such as dummy messages, mix public keys updates transmitted between mix nodes, etc.*
FPT_ITT.3	TSF data integrity monitoring	TSF data integrity monitoring requires that the TSF data transmitted between separate parts of the TOE is monitored for identified integrity errors.
FTP_TRP.1	Trusted path	

Table 8. Functional Components in the Multiple Mix Protection Profile (part 2)

Threat Label	Description
O.Anonymity	The TOE shall provide for an anonymous message delivery service; that is, the recipient of a message shall not be able to know the origin of the message, unless the author expressly inserts this information in the message body.
O.Untraceability	The TOE shall provide for an untraceable message delivery service; this means that, taken any message transiting through the system at any time, it shall not be possible to obtain enough information to link its origin and destination users.
T.ContentDisclosure	An attacker might intercept transiting messages between parts of the TOE and read their content, thus disclosing it, together with any related information. *This is a threat not only to the operation of the TOE (as discussed in [12]), but also for the user, whose communications might be traced. In particular, this threat relates to messages transiting from the user client to a node on the network and refers to both the original message content (written by the user), and also to the routing information and other auxiliary information carried by the message.*
T.EndPointTraffic Analysis	An attacker might intercept transiting messages between parts of the TOE (user client and mix node), and use the related information to perform traffic analysis on a user. *This threat relates to the concepts of sender anonymity and receiver anonymity. As viewed traditionally, main goal of the mix network is to hide the relation between receiver and sender of a message (this property also known as sender/receiver anonymity). However, once a suspect on a possible communication between two users is established, it may be possible to monitor the end points of message chains for a statistical correlation between transmission and reception times, especially if the traffic on the network is low, the users few, and the per-user traffic low. A similar discussion, related to Web transactions, may be found in [20].*
T.KeyForgery	An attacker might generate forged keys, simulating the activity of a given mix, distribute them, and make the user employ them to encrypt message in the belief that such messages are only readable by the replaced mix. *This is a threat to the originating user, who will send messages readable to an attacker, and might not be warned about it. A trust scheme (implemented for example by a certification authority) is required to counter this threat.*

Table 9. Threats in the User-Oriented Mix Protection Profile (part 1)

Threat Label	Description
T.Misuse	The user might install, configure or use the TOE interaction functions in an insecure manner, hence compromising the expected security properties offered by the TOE. *This threat is particularly relevant when considering the "human" element when this is the user, because the user is not expected to have as deep a knowledge about the TOE functions and about the security concerns as, for example, a system administrator, who represents the human element in the case of an administered mix node.*
T.OneStepPath	A mix may gain information linking origin and destination if the path from the origin user to the destination user contains only one mix.
T.UntrustworthyMix	Some mix(es) in the network may be compromised and hold, process and/or disclose information useful to trace, and/or reveal the content of, communications.
TE.MixConspiracy	Some mixes in the network may be compromised and share information useful to trace, and/or reveal the content of, communications. *This threat represents an extension to the T.UntrustworthyMix threat, in that it introduces the concept of information sharing between parts of the TOE.*
TE.PartialNetwork Block	An attacker might block the connection between parts of the TOE and the user. *This is a typical DoS attack, where part or the entire TOE is rendered unusable.*
TE.Redirection	An attacker might redirect the connections between parts of the TOE and act as to replace that part seamlessly, thus effectively acting as a compromised mix subset.

Table 9. Threats in the User-Oriented Mix Protection Profile (part 2)

Assumption Label	Description
A.SecurityGoals	The TOE is assumed to be used to achieve unlinkable and anonymous or pseudonymous communication. Other security properties, as unobservability of TOE use are not contemplated.
A.LogicalSec	The TOE will perform as long as the user takes care of securing the logical access to their computing environment. *This assumption requires some explanatory text. As logically securing mainstream operating systems and environments, especially when networked, is close to impossible*[2]*, the assumption should be taken rather loosely, provided that if the risk analysis leads to the conclusion that an attack on the user's workstation is likely, then the user should adopt a safer operating environment.*
A.OS	The single parts of the TOE run on operating system platforms that are assumed to be trusted and not to expose privacy related information belonging to the TOE.
A.PhysSec	The TOE will perform as long as the users take care of securing their physical access to the message traffic handled by the TOE. *This is a point that cannot be over-stressed; an insecure physical user location may be easily exploited against the user who mistakenly believes that his or her communications are unobserved.*
A.Minimal Connectivity	The TOE might not be able to reach its goal if an attacker is able to block all access points of the user to the mix network.
A.MinimalTrust	The TOE might not be able to reach its goal if all nodes (mixes) of the network are subverted.
A.OpenEnvironment	The mix network works in an open networked environment.
A.UnreliableNetwork	The connecting network might not be reliable on correctly delivering messages between parts of the TOE. Specifically, messages may be lost, altered or truncated accidentally. *The TOE is however not required to provide reliable service. A high degree of reliability may be achieved by sending multiple copies of a message through different paths.*
A.UserCooperation	Users cooperate actively at the enforcement of the security policy of the TOE. *Users are trusted to use in a correct manner the services made available by the TOE to reach their anonymity goals.*

Table 10. Assumptions in the User-Oriented Mix Protection Profile

Security Objective Label	**Description**
SO.Adequate Documentation	The TOE shall provide the user with adequate, readable documentation on the correct use of the security functions.
SO.Anonymity	The TOE shall accept and process messages without requiring that the processed data may be in any way linked to the origin user.
SO.ConcealMessage Content	The TOE shall enforce that the content of all messages transiting on the network be inaccessible to all third parties, in whatever point of the network the messages are intercepted.
SO.CounterTraffic Analysis	The TOE shall be constructed as to counter traffic analysis techniques specifically aimed at analyzing the communications between user client software and the mix network.
SO.DivideSecurity Information	The TOE shall be constructed as to provide the user the ability, and enforce the correct use of such ability, of determining the allocation of unlinkability-relevant data among different parts of the TOE.
SO.DivideSecurity Processing	The TOE shall provide to the user the ability, and enforce the correct use of such ability, of freely choosing a combination of mix nodes among which to allocate the processing activities achieving unlinkability.
SO.EnforceProper Use	The TOE (and especially the user interface part of the TOE) shall enforce the proper and secure use of the security functions of the TOE. *For example, require secure pass phrases, encryption, and minimum message chain length.*
SO.EnforceTrust Distribution	The TOE shall be constructed to enforce the user's choice of information and processing distribution.

Table 11. Security Objectives in the User-Oriented Mix Protection Profile (part 1)

Security Objective Label	**Description**
SO.Identity	The TOE shall uniquely identify the single mix nodes and users and provide means to transmit data to a specific mix while preserving the confidentiality of such data.
SO.KeyTrust Assurance	The TOE shall provide the user the ability, and enforce the correct use of such ability, of validating any public key used for encryption purposes against some trusted mechanism, to gain confidence that the communicating partner is actually who he claims to be.
SO.MinimizeSecurity Information	The TOE shall be constructed as to minimize the use, distribution and availability time frame of information impacting unlinkability.
SO.Untraceability	The TOE shall also ensure that no subject (user, administrator, threat agent) has the possibility to gain sufficient information as to track back the origin of a message.
SOE.Antagonistic Management	The TOE shall be independently and antagonistically managed. *The main problem with this security objective to be fulfilled by the environment is that it is nearly impossible to enforce it without some form of post-deployment assurance evaluation control and maintenance.*
SOE.Distributed Network	The TOE shall rely on a topologically distributed network. *This is required to maximize the resources that an attacker must deploy in the attempt to "cut off" part of the network from the rest. Apart from requiring specific design choices, this requirement can only be met by implementing a sound collective administration policy, and by providing means to assure the users of the effects of such a policy.*

Table 11. Security Objectives in the User-Oriented Mix Protection Profile (part 2)

	O.Anonymity	O.Untraceability	T.ContentDisclosure	T.EndPointTrafficAnalysis	T.KeyForgery	T.Misuse	T.OneStepPath	T.UntrustworthyMix	TE.MixConspiracy	TE.PartialNetworkBlock	TE.Redirection
SO.AdequateDocumentation						*					
SO.Anonymity	*										
SO.ConcealMessageContent			*								
SO.CounterTrafficAnalysis				*							
SO.DivideSecurityInformation							*	*			
SO.DivideSecurityProcessing							*	*			
SO.EnforceProperUse						*					
SO.EnforceTrustDistribution						*	*	*	*		
SO.Identity								*			*
SO.KeyTrustAssurance					*						
SO.MinimizeSecurityInformation								*			
SO.Untraceability		*									
SOE.AntagonisticManagement									*		
SOE.DistributedNetwork										*	*

Table 12. Security Objectives to Threats and Organizational Policies mapping

Sort name	Unique name	New?	Component Description and Comments
FCS_CKM.1	Crypto-graphic key generation		Cryptographic key generation requires cryptographic keys to be generated in accordance with a specified algorithm and key sizes that can be based on an assigned standard.
FCS_CKM.2	Crypto-graphic key distribution		Cryptographic key distribution requires cryptographic keys to be distributed in accordance with a specified distribution method that can be based on an assigned standard.
FCS_CKM.4	Crypto-graphic key destruction		Cryptographic key destruction requires cryptographic keys to be destroyed in accordance with a specified destruction method that can be based on an assigned standard.
FCS_COP.1	Crypto-graphic operation		Cryptographic operation requires a cryptographic operation to be performed in accordance with a specified algorithm and with a cryptographic key of specified sizes. The specified algorithm and cryptographic key sizes can be based on an assigned standard.
FDP_ACC.2	Complete access control		Complete access control requires that each identified access control SFP cover all operations on subjects and objects covered by that SFP. It further requires that all objects and operations with the TSC are covered by at least one identified access control SFP. *This access control policy, which is composed of this component and the following, states that:* – *All data produced by subjects covered by the SFP must obey the policy's requirements;* – *Data produced by subjects covered by the SFP must be explicitly addressed to some subject;* – *Data explicitly addressed to some subject must be unreadable by all other subjects;* – *Data produced by a subject may be read by the same subject that originated it.*
FDP_ACF.1	Security attribute based access control		Complete access control requires that each identified access control SFP cover all operations on subjects and objects covered by that SFP. It further requires that all objects and operations with the TSC are covered by at least one identified access control SFP.
FDP_IFC.1	Subset information flow control (CCE)		Subset information flow control requires that each identified information flow control SFP be in place for a subset of the possible operations on a subset of information flows in the TOE. *The CCE (Covert Channel Elimination) SFP, stated in this component and the following, requires the TOE to deploy techniques to eliminate covert channels by which an attacker may gain information about the use of the system by some user, especially with regards to traffic analysis information. (The specific technique to adopt is not specified.)*
FDP_IFF.4	Partial elimination of illicit information flows		Partial elimination of illicit information flows requires the SFP to cover the elimination of some (but not necessarily all) illicit information flows.

Table 13. Functional Components in the User-Oriented Protection Profile for unobservable message delivery using mix networks (part 1)

Sort name	**Unique name**	**New?**	**Component Description and Comments**
FDP_IRC.2	Full information retention control	Yes	Full information retention control requires that the TSF ensure that any copy of all objects in the TSC is deleted when not more strictly necessary for the operation of the TOE, and to identify and define the activities for which the object is required. *This component is used to state a minimization of access to information policy, which we tried to state using the stock CC components with access control requirements. However stating such a policy by means of access control is not satisfying, in that it represents a considerable extension to the intended use of the components, which are, as the name suggests, to be used to state information objects access policies in the traditional sense and do not lend themselves to other applications.* *For this reason this new component was developed and used in this PP.*
FDP_ITT.1	Basic internal transfer protection		Basic internal transfer protection requires that user data be protected when transmitted between parts of the TOE.
FDP_RIP.2	Full residual information protection		Full residual information protection requires that the TSF ensure that any residual information content of any resources is unavailable to all objects upon the resource's allocation or deallocation.
FIA_ATD.1	User attribute definition		User attribute definition, allows user security attributes for each user to be maintained individually.
FIA_UID.1	Timing of identification		Timing of identification, allows users to perform certain actions before being identified by the TSF.
FMT_MSA.1	Management of security attributes		Management of security attributes allows authorized users (roles) to manage the specified security attributes.
FMT_MSA.2	Secure security attributes		Secure security attributes ensures that values assigned to security attributes are valid with respect to the secure state.
FMT_MSA.3	Static attribute initialisation		Static attribute initialisation ensures that the default values of security attributes are appropriately either permissive or restrictive in nature.
FMT_SMR.1	Security roles		Security roles specifies the roles with respect to security that the TSF recognizes.
FPR_ANO.2	Anonymity without soliciting information		This component makes sure that the TOE does not request identification information regarding the origin and destination of messages it handles, and that nobody may gain information linking a data object (message) to users.

Table 13. Functional Components in the User-Oriented Protection Profile for unobservable message delivery using mix networks (part 2)

Sort name	**Unique name**	**New?**	**Component Description and Comments**
FPR_TRD.2	Allocation of information assets	Yes	Allocation of information assets requires that the TSF ensure that selected information impacting privacy be allocated among different parts of the TOE in such a way that in no state a single administrative domain will be able to access such information. *This component, and the following one, is needed to implement a trust distribution mechanism, which by the sole use of stock CC components was stated using the FPR_UNO.2 "Allocation of information impacting observability". However, the fact that it refers specifically to "unobservability" has impeded its use for other security properties. Additionally, in the initial version of the PP, which used the stock CC components, the FDP_ACC.2 "Complete access control", and FDP_ACF.1 "Security attribute based access control" were used to implement a mandatory access control policy in the TOE, which would require data:* – *To be explicitly addressed* – *To be not accessible by any subject except the intended addressee.* *However, using access control requirements to state requirements on the distribution of information resulted in stating unclear and ineffective requirements, since the structure of the CC access control components derives from experience in implementing standard access control policies, and does not lend itself well to the requirements needed for the mix.* *Therefore, in the new PP the more general FPR_TRD "Distribution of trust" family replaces all of the cited stock CC requirements components.* *This component divides the TOE (the mix network) in multiple administrative domains (a single mix node), as described in section 9.3.* *A more complete explanation of how this family enhances the PP can be found in section 9.3.3.*
FPR_TRD.3	Allocation of processing activities	Yes	FPR_TRD.3 Allocation of processing activities requires that the TSF ensure that selected processing activities impacting privacy be executed on different parts of the TOE in such a way that no single administrative domain will be able to make use of information gathered from the processing activity.
FPR_UNL.2	Unlinkability of users	Yes	Unlinkability of users requires that users and/or subjects are unable to determine whether two users are referenced to by the same object, subject or operation, or are linked in some other manner. *Originally the FPR_UNL.1 "Unlinkability" component was used to state requirements on the intended purpose of the mix network, i.e. to provide for unlinkable communication between partners. However, the fact that the CC unlinkability component is expressly limited to "unlinkability of operations" has made it difficult to use such a component in a more general way. For this reason it was replaced by the new, more general, FPR_UNL.2 "Unlinkability of users" component.*

Table 13. Functional Components in the User-Oriented Protection Profile for unobservable message delivery using mix networks (part 3)

The correspondence table between components and Objectives follows (Table 14). The tables provided in this section allow the reader to trace a single ECITS/CC component selected for inclusion in the PP to a specific threat or policy the TOE must counter or satisfy. Security Objectives that are not "covered" by any component must be addressed either by Assurance requirements, or by additional requirements on the environment, which are however not relevant at this point, and are here omitted.

	SO.AdequateDocumentation	SO.Anonymity	SO.ConcealMessageContent	SO.CounterTrafficAnalysis	SO.DivideSecurityInformation	SO.DivideSecurityProcessing	SO.EnforceProperUse	SO.EnforceTrustDistribution	SO.Identity	SO.KeyTrustAssurance	SO.MinimizeSecurityInformation	SO.Untraceability	SOE.AntagonisticManagement	SOE.DistributedNetwork
FCS_CKM.1									*					
FCS_CKM.2										*				
FCS_CKM.4										*				
FCS_COP.1			*											
FDP_ACC.2								*	*					
FDP_ACF.1								*	*					
FDP_IFC.1				*										
FDP_IFF.4				*										
FDP_IRC.2											*			
FDP_ITT.1			*											
FDP_RIP.2											*			
FIA_ATD.1									*					
FIA_UID.1					*	*								
FMT_MSA.1					*	*								
FMT_MSA.2					*	*	*							
FMT_MSA.3					*	*	*							
FMT_SMR.1					*	*								
FPR_ANO.2		*												
FPR_TRD.2					*			*					*	
FPR_TRD.3						*		*					*	
FPR_UNL.2												*		

Table 14. Functional components to Security Objectives mapping

The selected EAL level for this PP is EAL 5. The high assurance level is selected to gain a high level of assurance that the TOE will be developed, delivered, and evaluated following rigorous commercial practices. A formal model of the TOE security policies must be provided and evaluated, and the system must be independently tested. EAL 5 is the lowest level providing assurance components that impose the aforementioned tasks.

8 The experiences gained

The process of writing PPs is supposed to be top-down. The author identifies a set of threats, devises a set of security objectives that should counter all the threats, and finally expresses these objectives through a set of formal requirements taken from the ECITS/CC catalog. This methodology has many advantages, the main one being that the development process of the PP is clean, and the formal demonstration of correspondence between the various threats, objectives and requirements is simple.

The problems arise when the PP author needs to express requirements for security objectives not covered by ECITS/CC components. During the development of the User-Oriented Protection Profile, three such issues were identified:

1. Requirements on the distribution of the TOE: although it may be viewed as a purely architectural requirement, it is worthy to note that many secure systems are based explicitly on a distributed system to perform the security relevant tasks. Mixes are an example, but also digital payment systems, etc. show such pattern.
2. Requirements on the policies requiring the minimization of knowledge: clearly information that has been disposed of cannot be disclosed. Deleting information as soon as it is not essential to the operation of the system anymore is thus always a safe practice.
3. Requirements on unlinkability properties to be enforced by the TOE: the statement of unlinkability of operations is possible through the stock ECITS/CC components, but not so for unlinkability of users, which is precisely what the mix network provides.

To solve the expressive deficiencies of the ECITS/CC a number of options may be considered, and the following three are worthwhile to mention:

1. Restate the security objective differently, (i.e. “fit” the objective to the requirements),
2. Try to force the criteria components to cover the objective (i.e. “fit” the requirements to the objective),
3. Develop new functional components.

The first two options show to be not viable in the long run. In fact, the first one breaks the top-down paradigm, and distorts the PP to state what is expressible by the criteria, necessarily avoiding all security issues which are not simply stateable by the ECITS/CC. The second option “overloads” the ECITS/CC components to express requirements for which they were not thought. This has many drawbacks; for one thing, it may simply be not always possible. Moreover, the requirements tend to become unclear, and ineffective, and the PP evaluation process becomes more complicated because of the non-straightforward use of the components.

The third option has undoubtedly many formal and theoretical advantages, and some drawbacks. On the one hand, the requirements may be stated in a simple fashion, and the top-down structure is preserved. On the other hand, while the ECITS/CC allow for expansion of the base requirements sets, one of their main advantages (i.e. mutual recognition) is not guaranteed for PPs that use such novel components.

The full discussion of the various problems encountered, and of how it was decided to write new components is too lengthy to be included here, but it can be said that each of the previous issues arose when trying to express specific objectives through the criteria, and an effort was made to approach the problem by using all three strategies [14]. In each case, the conclusion was found that the technically best way to proceed was that of developing new components.

The decision might have been different in the situation of a concrete evaluation. There resource constraints (getting an evaluation through without spending too much time discussing novel approaches) and easier mutual recognition (therefore staying with the standard set of components) might have got priority. However, with respect to improving the ECITS/CC two new families and one largely revised family are proposed in the next chapter.

9 Proposals for new and revised functional families

Three new functional families were devised, in a general enough formulation, and in a suitable format to be included in the ECITS/CC set. The new families are summarized inTable 15. Each family is discussed in a separate section below. The formal statement of the three families, which follows the typographical, layout and content standards of the ECITS/CC, can be found in the Annexes (Chapter 12).

The components were proposed to solve precise problems we incurred in while using the ECITS/CC to state requirements for mix networks, but are devised to be as reusable and general as possible.

Label	Name	Purpose
FDP_IRC	Information retention control	Limit the accumulation of non-essential information.
FPR_UNL	Unlinkability	Extend the current unlinkability requirements.
FPR_TRD	Distribution of trust	Allow the user to allocate information and processing activities.

Table 15. Proposed new and revised functional families

9.1 Information retention control (FDP_IRC)

The "Information retention control" family addresses a basic need in secure information processing and storage applications, which however appears not to be covered by the ECITS/CC: the need for secure management of data no more needed by the TOE to perform its operation, but still stored in the TOE. The traditional view of IT systems as data storage systems induced naturally into thinking that once entered, data would be seldom deleted from the system, and if so, mainly because of storage exhaustion problems.

But in a multilateral or high security environment it is important to minimize the replication, and temporal time frame in which information is contained in the system. Also, users might want their IT products to avoid retaining data that they consider exploitable by third parties, or threatening their privacy. In this case, such a requirement can help users to gain confidence that the product is secure, as far as it deletes every copy of the data when not needed anymore.

The FDP_RIP "Residual information protection" family addresses one side of this problem[3], but an explicit requirement on the management of no longer needed data is missing.

Of course competing requirements may arise, as data may be needed by the system for more activities over a long period of time. Possible solutions to this problem are:

- Better protecting the information objects stored in the TOE from access,
- Re-requesting the protected information from the user each time it is needed.

Overview of the family Information retention control ensures, that information no longer necessary for the operation of the TOE is deleted by the TOE. Components of this family require the PP author to identify TOE activities and objects required for those activities, and not to be kept in the TOE, and the TOE to keep track of such stored objects, and to delete on-line and off-line copies of unnecessary information objects.

The suggested class for this family is class FDP "*User Data Protection*", since the main purpose of this family is the protection of user data while in the TOE.

This family sets only requirements on information objects requested for specific activities in the TOE operation, and not on general data gathering. The policies which control the collection, storage, processing, disclosure and elimination of general user data stored on the TOE must be detailed elsewhere, and are domain of the environmental objectives and organizational policies, not of the PP.

Components belonging to this family could be used, for example, when the TOE needs some information from the user, or generates information, which might be easily mismanaged or misused in case of a malicious or inadvertent use or administration of the TOE. This category includes, for example:

[3] Namely, the elimination from the TOE of all traces left behind by objects upon deallocation of resources used to store or manipulate them.

- Connecting IP numbers on anonymizing proxy servers;
- One-time cryptographic keys, which if eventually disclosed could allow the decryption of intercepted and stored communications or files;
- TOE usage histories, as interactive command line shell histories, or information presentation tools cache files (i.e. WWW browser caches), which, while useful to the user and TOE during specific activities, could be used to track user or TOE actions, if preserved across sessions;
- Any information for which security considerations (both of the TOE and of the user) suggest not to keep on the TOE, if not strictly necessary.

When more than one activity requires the presence of a protected object, all activities, which refer to the required object must end before deleting it.

Components The family has two hierarchical components:

FDP_IRC.1 Subset information retention control requires that the TSF ensure that any copy of *a defined subset of* objects in the TSC is deleted when no longer strictly necessary for the operation of the TOE, and to identify and define the activities for which the object is required.
FDP_IRC.2 Full information retention control requires them same but regarding to *all* objects in the TSC.

FDP_IRC.1 Subset information retention control This component requires that, for a subset of the objects in the TOE, the TSF will ensure that the objects will be deleted from the TOE when no longer required for some specific action.

The formal description of the component is available in section 12.1. The PP/ST author should specify the list of objects subject to information retention control. He should also specify the list of activities which require specific objects to be stored in the TOE, and whose termination requires the TOE to delete the no more required objects.

FDP_IRC.2 Full information retention control This component requires that, for all objects in the TOE, the TSF will ensure that the objects will be deleted from the TOE when no longer required for some specific action. In other words, *every* object used by the TOE must be tracked for its necessity, and if not more strictly required, deleted. Therefore this component is hierarchical to FDP_IRC.1.

The assignment can be limited to specifying the list of activities which require specific objects to be stored in the TOE, and whose conclusion requires the TOE to delete the no more required objects.

9.2 Unlinkability (FPR_UNL)

The general model of entities as set up in the ECITS/CC (cf. 4.1) allows specifying various kinds of security requirements, including privacy-related requirements. For example an unlinkability of operations requirement would impose a

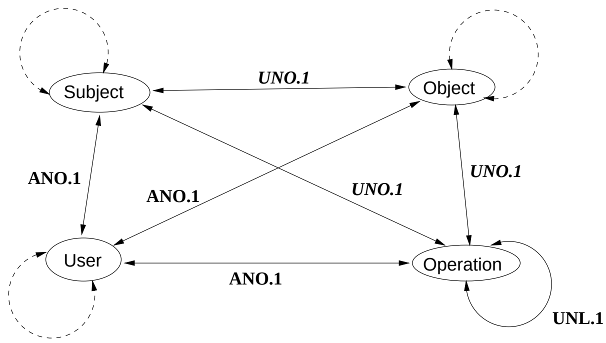

Fig. 3. Unlinkability properties covered (solid arrows) and not covered (dashed arrows) by existing components

constraint on the relationship between operations in the TSC relating them to a particular user.

However, the full expressive potential of this model is not described by the standard ECITS/CC components. Figure 3 shows the current situation: The solid arrows indicate a relationship, which is covered by a particular ECITS/CC component, and the dashed arrows indicate that the link they represent is not expressible using the current ECITS/CC privacy components. With regard to unlinkability, the ECITS/CC provide the FPR_UNL.1 component that provides unlinkability of operations (cf. 4.2.3). Its only functional element reads:

> **FPR_UNL.1.1** The TSF shall ensure that [assignment: *set of users and/or subjects*] are unable to determine whether [assignment: *list of operations*] [selection: *were caused by the same user, are related as follows* [assignment: list of relations]].

Although useful, this family does not cover at least one case, which is of primary importance for mixes: the unlinkability of users, in relation to a specific data object (the mail message). This kind of property is also hard to express through the other families: one could try using the unobservability (FPR_UNO) family, which is however not adequate because the action itself of transmitting a message is not hidden by the mix system. The mix hides only the relation between users, and between email and user.

In conclusion an enhancement of the unlinkability family is necessary to augment the expressiveness of the ECITS/CC to include also the mentioned cases.

Overview of the family The aim of the unlinkability family is still to ensure that selected entities may refer each another without others being able to observe these references (cf. 4.2.3); the change is that it now applies not only to users operations, but also to subjects and objects.

The components share a common structure and provide the PP author with the possibility of tailoring the following:

1. The users and subjects from which the information should be hidden.
2. A list of specific entities that the requirement protects.
3. A selection or an assignment of a list of relationships to hide.

Components The family consists of four sibling components:

FPR_UNL.1 Unlinkability of operations requires that users and/or subjects are unable to determine whether the same user caused certain specific operations in the system, or whether operations are related in some other manner. This component ensures that users cannot link different operations in the system and thereby obtain information.
FPR_UNL.2 Unlinkability of users requires that users and/or subjects are unable to determine whether two users are referenced to by the same object, subject or operation, or are linked in some other manner.
This component ensures that users cannot link different users of the system and thereby obtain information on the communication patterns and relationships between users.
FPR_UNL.3 Unlinkability of subjects requires that users and/or subjects are unable to determine whether two subjects are referenced to by the same object, user or operation, or are linked in some other manner.
This component ensures that users cannot link different subjects in the system and thereby obtain information on the usage and operation patterns of the subjects.
FPR_UNL.4 Unlinkability of objects requires that users and/or subjects are unable to determine whether two objects are associated to the same user, subject or operation, or are linked in some other manner.
This component ensures that users cannot link different objects in the system and thereby obtain information on the usage patterns of objects.

9.3 Distribution of trust (FPR_TRD)

Among the current families in the privacy class of the ECITS/CC no provision is made to address privacy requirements related to the distribution of trust among parts of the TOE, except in the FPR_UNO.2 component; the new functional family is therefore proposed to be integrated into the FPR class.

Trust may be defined, not only in an IT setting, as *"Assured resting of the mind on the integrity, veracity, justice, friendship, or other sound principle, of another person; confidence; reliance."* [23]. In a more restrictive definition, one may define it as *"confidence on the integrity of another person or organization in the managing of an asset given to him, her or it"*. In this context, trust division may be described as the process of allocating assets among different trustees with the aim of minimizing the damage, which one might suffer if one of the trustees betrays the trust given.

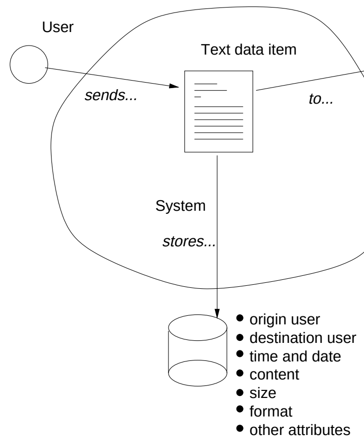

Fig. 4. Hidden activities and information objects involved in sending a data object through a system

Clearly in IT the main asset is *information*, and the accidental or intentional loss or mismanagement of it may result in great damages for the owners or beneficiaries of it. Data may be either supplied directly to an information system, as inputted files, documents, personal information, or they may be derived from interaction with the system, such as data regarding on-line time and login times of a user, requests and destination of email deliveries and WWW accesses, or called telephone numbers; often the collection of this kind of information is not clearly stated in the (contract) terms which bind user and operator of a system. Figure 4 shows the hidden information processed and possibly stored in a system, which provides textual data transmission capabilities to users when such an operation is initiated.

Another related observation is that the processing itself produces information, whose existence or content may not even be known to the user that requested the processing activity to be initiated. For example, in large WWW sites which employ distributed, redundant, servers requests are redirected to one of the servers in a pool, and such mechanism, and also the identity of the server which actually executes the request is not visible to the end user, neither is the server choice known.

The proposed *"Distribution of trust"* family addresses both aspects of the trust issue, i.e. the distribution of information, and the distribution of processing activities, which may produce privacy-relevant information themselves.

Overview of the family This family describes specific functions that can be used to allocate information and processing activities on the TOE with the objective of protecting the privacy of users of the system. To allow such allocation, the concept of *"Administrative Domain"* (AD) is introduced to indicate a part of the TOE whose security functions are accessible and usable to access data by a single subject (system user, administrator ...) without requesting any additional authorization or performing additional authentication procedures.

The AD is a formalization of the concept of the more intuitive "part of the TOE", which is also used in the statement of the FPR_UNO.2 component. Moreover, it specifies that administrators of an AD may not access other ADs without gaining rightful permission. In fact, allocating information on different "parts of the TOE" is not very useful, if the different parts are accessible by the same administration. If all the parts are administered by the same user or organization, a subverted administrator or an attacker gaining administrator privileges may as well access such information even if it is distributed. Instead, it is necessary to provide for independent administration and separate access domains for different parts of the TOE; this means that an administrator of one part will not be able to access as such also other parts of the TOE.

As an example, consider a monolithic TOE (i.e. a UNIX operating system environment), where only one administrative domain exists, and the administrator may access the security functions of the whole TOE. As a result, if users store both their sensitive data, even in an encrypted form, and their private keys on the same system, the administrator (or an attacker gaining administrator privileges) will be able to access the data.

To avoid this problem, the TSF could be designed to allocate data and keys on different, independently administered systems, and to require that the decryption be done on a third system when the owner needs to access it. This obviously raises the common chicken and egg problem of whether the system where the cryptographic functions take place is trusted or not. Many solutions can be applied in this case, e.g.:

1. Performing a two-phase en/decryption in separate administrative domains (which is, in essence, what the mix system does),
2. Personally administering the system where cryptographic functions take place (for example, a smartcard with cryptographic capabilities, which stores the keys and communicates with the outside only with the input and output of cryptographic algorithms; the card is always carried by the owner of the data, which trusts the issuer of the card, or a certificate regarding the card[4].)

Components The family is structured in three components, one of which is a base component defining the concept of administrative domain, while the other two express the requirements on information and operations allocation:

> *FPR_TRD.1 Administrative domains* requires that the TOE be divided in distinct administrative domains (AD), with separate authentication and access control procedures; administrators of one administrative domain may not access other ADs.

[4] Of course a secure administration would also require secure input (e.g. keyboard) and output (e.g. display) facility for the user.

FPR_ TRD.2 Allocation of information assets requires that the TSF ensure that selected information impacting privacy be allocated among different parts of the TOE in such a way that in no state a single administrative domain will be able to access such information.
FPR_ TRD.3 Allocation of processing activities requires that the TSF ensure that selected processing activities impacting privacy be executed on different parts of the TOE in such a way that no single administrative domain will be able to make use of information gathered from the processing activity.

The derivate components (FPR_TRD.2 and FPR_TRD.3) let the PP author tailor the following:

1. A list of objects or operations which should be subject to allocation in different ADs,
2. In the case of objects, the form of allocation (e.g. distribution, encryption ...),
3. A set of conditions that should always be maintained by the TOE with regard to assets allocation.

The formal component descriptions are available in section 12.3.

The effect of using the new components In the previous chapters, we stated that the introduction of the new privacy-oriented components in the User-Oriented Mix PP greatly simplified the statement of the requirements and enhanced their effectiveness. To support this assertion, we now show in detail how the new components perform. To avoid lengthening the paper excessively we will limit the example to only one of the new functional families (FPR_TRD "Distribution of Trust").

The introduction of the new components has a twofold advantage. First of all, it allows requirements to be specified in a more clear and simple manner compared to using the stock components, which had to be overloaded to express certain requirements for which they were not intended. Secondarily, it also allows expressing more complete and precise requirements, and reduces the number of unmet Security Objectives.

Table 16 shows the subset of security objectives in which the new FPR_TRD family is used. For each Security Objective, the table lists the stock functional components that were used in the first version of the PP (second column), and the components used in the final version (third column).

The new components do not only replace some of the old ones, but also provide for a better coverage of the security objectives stated in the PP. The following list shows this in detail for every security objective:

- SO.DivideSecurityInformation "The TOE shall be constructed as to allow the user the ability, and enforce the correct use of such ability, the allocation of unlinkability-relevant data among different parts of the TOE."
  Before the introduction of the new families, this objective was reached by adopting a set of three requirements. Essentially, an access control policy

Security Objective	**Initial PP**	**Final PP**
SO.Divide Security Information	FDP_ACC.2 “Complete access control (MUDAC)” FDP_ACF.1 “Security attribute based access control (MUDAC)” FMT_MSA.1 “Management of security attributes” FMT_MSA.2 “Secure security attributes” FMT_MSA.3 “Static attribute initialisation” FPR_UNO.2 “Allocation of information impacting unobservability”	FMT_MSA.1 “Management of security attributes” FMT_MSA.2 “Secure security attributes” FMT_MSA.3 “Static attribute initialisation” FPR_TRD.2 “Allocation of information assets”
SO.Divide Security Processing	FMT_MSA.1 “Management of security attributes” FMT_MSA.2 “Secure security attributes” FMT_MSA.3 “Static attribute initialisation”	FMT_MSA.1 “Management of security attributes” FMT_MSA.2 “Secure security attributes” FMT_MSA.3 “Static attribute initialisation” FPR_TRD.3 “Allocation of processing activities”
SO.EnforceTrust Distribution	FDP_ACC.2 “Complete access control (MUDAC)” FDP_ACF.1 “Security attribute based access control (MUDAC)”	FDP_ACC.2 “Complete access control (MUDAC)” FDP_ACF.1 “Security attribute based access control (MUDAC)” FPR_TRD.2 “Allocation of information assets” FPR_TRD.3 “Allocation of processing activities”
SOE.Antagonistic Management	Previously no component available to cover this objective	FPR_TRD.2 “Allocation of information assets” FPR_TRD.3 “Allocation of processing activities”

Table 16. How the FPR_TRD family helps to fulfill Security Objectives

would control the enforcement part of the requirement, while the security attribute management components would allow the user to divide the allocation of security-relevant information. Finally, the "Allocation of information impacting unobservability" component was used in an "overloaded" manner, which proved to be ineffective. Thus, the initial PP addressed this Security Objective by using the following components:

- FPR_UNO.2 "Allocation of information impacting unobservability"
  This is the only component in the CC/ECITS that expressly provides for allocation of information. However, the fact that it refers specifically to "unobservability" causes problems to its use for other security properties. The "trick" for overloading the stock component was that of requiring the operation of transmitting a message between users to be unobservable. However, this results in an ambiguous requirement because nothing can be said about the *link* between communicating partners, which a mix network also aims at hiding (Unlinkability).
- FDP_ACC.2 "Complete access control (MUDAC)", and FDP_ACF.1 "Security attribute based access control (MUDAC)"
  These two components were introduced into the initial PP to implement a mandatory access control policy. This policy requires data to be explicitly addressed and access to be strictly controlled and limited to the intended recipient. The components remain in the new PP to enforce the SO.EnforceTrustDistribution and SO.Identity objectives, but are superseded by the FPR_TRD.2 "Allocation of information assets" component with regard to the SO.DivideSecurityInformation objective.
  In this case, the access control requirements allow the PP author to define requirements on which subjects may access the information that flows through the mix network. However, they fail completely at specifying requirements on how such information flow must be structured to achieve unlinkability and unobservability (the distributed nature of message processing in the mix network).

In the final version, the division of trust component takes the place of both the access control components and the allocation of unobservability information component.

– SO.DivideSecurityProcessing "The TOE shall provide to the user the ability, and enforce the correct use of such ability, of freely choosing a combination of mix nodes among which to allocate the processing activities achieving unlinkability."
  In this case the objective was not fully satisfied in the initial version of the PP, because the CC/ECITS do not provide a functional component for allocating *processing activities* in different parts of the TOE.
  This previously not satisfied objective can now be fully covered by using one of the new components. The new FPR_TRD.3 "Allocation of processing activities" component provides for distribution of processing among different, independently administered, parts of the TOE, while the ability for the user to specify some of the security attributes (which is how routing information is considered in this PP) allows to actually make use of distributed processing.

- SO.EnforceTrustDistribution "The TOE shall be constructed to enforce the user's choice of information and processing distribution."

  This requirement was only partially covered in the initial PP, because the access control requirements do not allow stating requirements on the TOE structure. Adding the FPR_TRD components complements the access control requirements and results in a fully covered objective.
- SOE.AntagonisticManagement "The TOE shall be independently and antagonistically managed."

  This objective that was not at all covered in the first version of the PP is now partially covered, as the TOE is now built to allow for independent administration, at least from a technical point of view. Obviously adequate environmental procedures and policies are still necessary for the correct operation of the TOE.

To ease analyzing the relationship between security objectives and functional components, Table 17 splits the objectives in atomic assertions and shows how each assertion is covered by one or more components.

Requirement name	**Single component statements**	**Satisfied by ...**
SO.Divide Security Information	TOE shall be distributed	Administrative domains FPR_TRD
SO.Divide Security Processing	User ability of choosing a distributed use pattern	Management of security attributes FMT_MSA
SO.EnforceTrust Distribution	Enforce users' choices	Mandatory access control FDP_ACF, FDP_ACC
	Construction of the TOE as to allow information and processing distribution	Administrative domains FPR_TRD
SOE.Antagonistic Management	Independent management	Administrative domains FPR_TRD

Table 17. Overview of coverage of the TOE distribution objective

Note that objectives and functional components do not match exactly, i.e. more than one component is necessary to meet a security objective, and a single component may address more than one objective. This is a common situation when both the objectives and the components state complex requirements with multiple, independent assertions.

As a final note one may observe that in the old PP, without the trust division components, the partial objectives marked in Table 17, as "construction of the TOE" and "TOE shall be distributed" were simply not covered.

10 Summary and conclusion

The experiences gained while writing the Protection Profiles include the following major issues:

1. In general the ECITS/CC provide much more flexibility than their predecessors. They also contain much better instruments to describe privacy friendly functionality. However as shown above, the ECITS/CC components do not offer a complete solution to all the issues which characterize privacy-related objectives.
2. The greatest challenges to the expressive capacity of the functional components appear in the Multiple Mix PP and in the User-Oriented Mix PP, where a point of multilateral security is raised (security of the TOE vs. security of the user).
3. For some applications, architectural choices and objectives (i.e. distributed vs. centralized system) influence the security properties of the system. This applies to mixes, but holds also for other "secure" applications, as digital money, information handling and storage, etc.
4. The probably most relevant evidence is that simply trying to force the application's requirements or the functional components to "fit" is not a sustainable solution, because it results in an unclear and ineffective requirements definition.
5. The proposed components aim at forming a useful start-up for enhancing future versions of the ECITS/CC, even when the respective part of the criteria becomes slightly longer. Privacy oriented functionality covers only a small part (ca. 10 percent) of the criteria, so there should be space for the improvements.
6. Especially in the area of communication the evaluation of service security becomes important for users. While the ECITS/CC provide some help for this further work is needed.

Annexes

References

1. British Standards Institution: Code of practice for information security management (BS 7799-1: 1999); Specification for information security management systems (BS 7799-2: 1999)
2. c:cure scheme; http://www.c-cure.org
3. Common Criteria Implementation Board: Common Criteria for IT Security Evaluation, V. 2.0, May 1998; http://csrc.nist.gov/cc
4. Common Criteria Implementation Board: Common Criteria for IT Security Evaluation, V. 2.1, August 1999; www.commoncriteria.org and http://csrc.nist.gov/cc
5. Common Criteria Project: List of Protection Profiles; http://csrc.nist.gov/cc/pp/pplist.htm

6. European Commission: IT Security Evaluation Criteria, V. 1.2; 1991-06-28; Office for Official Publications of the EC; also www.itsec.gov.uk/docs/pdfs/formal/ITSEC.PDF
7. Untraceable electronic mail, return addresses, and digital pseudonyms. Communications of the ACM, 1981, Vol. 24, No. 2, pp. 84-88
8. Chris Corbett: ITSEC in Operation - an Evaluation Experience, Proc. 4th Annual Canadian Computer Security Conference, May 1992, Ottawa, Canada, pp. 439-460
9. Lance Cottrell: Mixmaster & Remailer Attacks; http://www.obscura.com/~loki/remailer/remailer-essay.html
10. Privacy Protection and Data Security task Force of the German Society for Informatics: Statement of Observations concerning the Information Technology Security Evaluation Criteria (ITSEC) V1.2; 24 February 1992, edited in Data Security Letter, No 32, April 1992
11. Giovanni Iachello: Single Mix Protection Profile, Revision 1.11, May 1999; http://www.iig.uni-freiburg.de/~giac
12. Giovanni Iachello: Protection Profile for an Unobservable Message Delivery Application using Mixes, Revision 1.7, June 1999; http://www.iig.uni-freiburg.de/~giac
13. Giovanni Iachello: User-Oriented Protection Profile for an Unobservable Message Delivery Application using Mix networks, Revision 2.4, June 1999; http://www.iig.uni-freiburg.de/~giac
14. Giovanni Iachello: IT Security Evaluation Criteria, and Advanced Technologies for Multilateral Security - The Mix Example; Tesi di Laurea; Universität Freiburg, Institut für Informatik und Gesellschaft, Abt. Telematik and Università degli Studi di Padova; June 1999; http://www.iig.uni-freiburg.de/~giac
15. ISO/IEC: Guidelines for the management of IT security (GMITS); Parts 1-5; Technical Report 13335 (part 5 still under development)
16. ISO/IEC: Evaluation Criteria for IT Security (ECITS), Parts 1-3; International Standard 15408;1999-12-16
17. Anja Jerichow, Jan Müller, Andreas Pfitzmann, Birgit Pfitzmann, Michael Waidner: Real-Time Mixes: A Bandwidth-Efficient Anonymity Protocol; IEEE Journal on Selected Areas in Communications 16/4 (May 1998) 495-509
18. Kai Rannenberg: Recent Development in Information Technology Security Evaluation - The Need for Evaluation Criteria for multilateral Security; in Richard Sizer, Louise Yngström, Henrik Kaspersen und Simone Fischer-Hübner: Security and Control of Information Technology in Society - Proceedings of the IFIP TC9/WG 9.6 Working Conference August 12-17, 1993, onboard M/S Ilich and ashore at St. Petersburg, Russia; North-Holland, Amsterdam 1994, pp. 113-128; ISBN 0-444-81831-6
19. Kai Rannenberg: What can IT Security Certification do for Multilateral Security? pp. 515-530 in Günter Müller, Kai Rannenberg: Multilateral Security in Communications - Technology, Infrastructure, Economy; Addison-Wesley-Longman, München, Reading (Massachusetts) ... 1999; ISBN 3-8273-1360-0
20. M. K. Reiter and A. D. Rubin. Crowds: Anonymity for web transactions. ACM Transactions on Information and System Security 1(1):66-92, November 1998.
21. Paul F. Syverson, David M. Goldschlag, Michael G. Reed: Anonymous connections and onion routing; in: Proceedings of the 1997 IEEE Symposium on Security and Privacy; IEEE Pres, Piscataway NJ
22. USA Department of Defense Trusted Computer System Evaluation Criteria; December 1985, DOD 5200.28-STD, http://www.radium.ncsc.mil/tpep/library/rainbow/5200.28-STD.html

23. Webster's Revised Unabridged Dictionary, 1913, ftp://ftp.dict.org/pub/dict/
24. The Freedom Network Architecture, Version 1.0, Zero-Knowledge Systems, Inc., http://www.zks.net/products/whitepapers.asp

11 Proposed criteria components

The three proposed families are included in a notation conformant with the prescriptions of the ECITS/CC. Specifically, this means that **bold facing** in components or in parts of components indicates an additional requirement compared with a hierarchical lower component.

11.1 Information retention control (FDP_IRC)

Family behaviour

This family addresses the need to ensure that information no longer necessary for the operation of the TOE is deleted by the TOE. Components of this family require the PP author to identify TOE activities and objects required for those activities, and not to be kept in the TOE, and the TOE to keep track of such stored objects, and to delete on-line and off-line copies of unnecessary information objects.

Component levelling

FDP_IRC.1 Subset information retention control requires that the TSF ensure that any copy of a defined subset of objects in the TSC is deleted when not more strictly necessary for the operation of the TOE, and to identify and define the activities for which the object is required.

FDP_IRC.2 Full information retention control requires that the TSF ensure that any copy of all objects in the TSC is deleted when not more strictly necessary for the operation of the TOE, and to identify and define the activities for which the object is required.

Management: FDP_IRC.1, FDP_IRC.2

There are no management activities foreseen for this component.

Audit: FDP_IRC.1, FDP_IRC.2

There are no events identified that should be auditable if FAU_GEN Security audit data generation is included in the PP/ST.

FDP_IRC.1 Subset information retention control

Hierarchical to: No other components

FDP_IRC.1.1 The TSF shall ensure that [assignment: *list of objects*] required for [assignment: *list of activities*] shall be eliminated immediately from the TOE upon termination of the activities for which they are required.

Dependencies: No dependencies.

FDP_IRC.2 Full information retention control

Hierarchical to: FDP_IRC.1

FDP_IRC.2.1 The TSF shall ensure that **all** objects required for [assignment: *list of activities*] shall be eliminated immediately from the TOE upon termination of the activities for which they are required.

Dependencies: No dependencies.

11.2 Unlinkability (FDP_UNL)

This family ensures that selected entities may be linked together without others being able to observe these links.

Component levelling

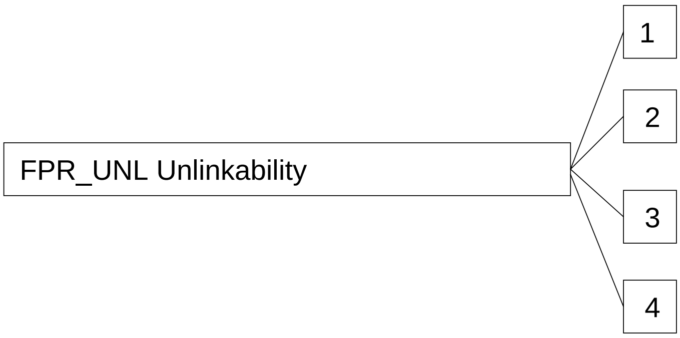

FPR_UNL.1 Unlinkability of operations requires that users and/or subjects are unable to determine whether the same user caused certain specific operations in the system, or are related in some other manner.

FPR_UNL.2 Unlinkability of users requires that users and/or subjects are unable to determine whether two users are referenced to by the same object, subject or operation, or are linked in some other manner.

FPR_UNL.3 Unlinkability of subjects requires that users and/or subjects are unable to determine whether two subjects are referenced to by the same object, user or operation, or are linked in some other manner.

FPR_UNL.4 Unlinkability of objects requires that users and/or subjects are unable to determine whether two objects are associated to the same user, subject or operation, or are linked in some other manner.

Management: FPR_UNL.1, FPR_UNL.2, FPR_UNL.3, FPR_UNL.4

The following actions could be considered for the management functions in FMT:

a) the management of the unlinkability function.

Audit: FPR_UNL.1, FPR_UNL.2, FPR_UNL.3, FPR_UNL.4

The following actions shall be auditable if FAU_GEN Security audit data generation is included in the PP / ST:

a) Minimal: The invocation of the unlinkability mechanism.

FPR_UNL.1 Unlinkability of operations

Hierarchical to: No other components

FPR_UNL.1.1 The TSF shall ensure that [assignment: *set of users and/or subjects*] are unable to determine whether [assignment: *list of operations*] [selection: *were caused by the same user, are related as follows* [assignment: *list of relations*]].

Dependencies: No dependencies.

FPR_UNL.2 Unlinkability of users

Hierarchical to: No other components

FPR_UNL.2.1 The TSF shall ensure that [assignment: *set of users and/or subjects*] are unable to determine whether [assignment: *list of users*] [selection: *are referenced by the same operation, are referenced by the same object, are referenced by the same subject, are related as follows* [assignment: *list of relations*]].

Dependencies: No dependencies.

FPR_UNL.3 Unlinkability of subjects

Hierarchical to: No other components

FPR_UNL.3.1 The TSF shall ensure that [assignment: *set of users and/or subjects*] are unable to determine whether [assignment: *list of subjects*] [selection: *act on behalf of the same user, are referenced by the same object, are referenced by the same operation, are related as follows* [assignment: *list of relations*]].

Dependencies: No dependencies.

FPR_UNL.4 Unlinkability of objects

Hierarchical to: No other components

FPR_UNL.4.1 The TSF shall ensure that [assignment: *set of users and/or subjects*] are unable to determine whether [assignment: *list of objects*] [selection: *are associated to the same user, are associated to the same subject, are associated to the same operation, are related as follows* [assignment: *list of relations*]].

Dependencies: No dependencies.

11.3 Distribution of trust (FPR_TRD)

This family addresses the need to ensure that privacy-relevant information referring to a user of a TOE is divided among different parts of the TOE, or stored in such a manner (as with encryption) to make it impossible that a part of the TOE under a single administrative domain is able to access such information.

Component Levelling

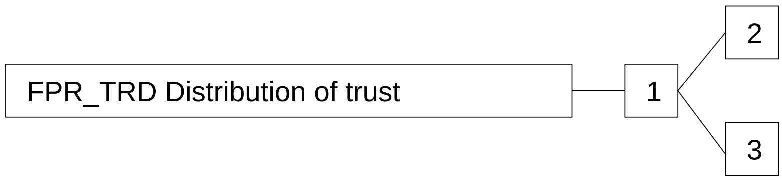

FPR_TRD.1 Administrative domains requires that the TOE be divided in distinct administrative domains (AD), with separate authentication and access control procedures; administrators of one administrative domain may not access other ADs.

FPR_TRD.2 Allocation of information assets requires that the TSF ensure that selected information impacting privacy be allocated among different parts of the TOE in such a way that in no state a single administrative domain will be able to access such information.

FPR_TRD.3 Allocation of processing activities requires that the TSF ensure that selected processing activities impacting privacy be executed on different parts of the TOE in such a way that no single administrative domain will be able to make use of information gathered from the processing activity.

Management: FPR_TRD.1

There are no management activities foreseen for this component.

Management: FPR_TRD.2

The following actions and definitions could be considered for the management functions in FMT:

1. The FMT_SMR.1 component could define a new security role “information owner” with regard to a specific data object or operation; this role represents the originator, and main user and beneficiary of such object or operation, and is the only subject or user allowed to specify distribution policies as security attributes for these entities;
2. An information owner could define default object security attributes;
3. An information owner could define and change security attributes on objects he or she owns.

Management: FPR_TRD.3

The following actions and definitions could be considered for the management functions in FMT:

1. The FMT_SMR component could define a new security role “information owner” with regard to a specific data object or operation; this role represents

the originator, and main user and beneficiary of such object or operation, and is the only subject or user allowed to specify distribution policies as security attributes for these entities;
2. An information owner could define default operation security attributes;
3. An information owner could define and change security attributes on operations it initiates.

Audit: FPR_TRD.1, FPR_TRD.2, FPR_TRD.3

There are no events identified that should be auditable if FAU_GEN Security audit data generation is included in the PP/ST.

FPR_TRD.1 Administrative domains

Hierarchical to: No other components

FPR_TRD.1.1 The TOE shall be divided in separate, independent, intercommunicating parts (administrative domains) governed by distinct access control and authentication configurations.

FPR_TRD.1.2 The distinct administrative domains of the TOE shall explicitly request access to data stored on other parts of the TOE to be granted access to it.

Dependencies: No dependencies.

FPR_TRD.2 Allocation of information assets

Hierarchical to: FPR_TRD.1

FPR_TRD.2.1 The TOE shall be divided in separate, independent, intercommunicating parts (administrative domains) governed by distinct access control and authentication configurations.

FPR_TRD.2.2 The distinct administrative domains of the TOE shall explicitly request access to data stored on other parts of the TOE to be granted access to it.

FPR_TRD.2.3 The TSF shall ensure that [assignment: *list of objects*] shall be stored [selection: *on different administrative domains of the TOE, in a form unreadable by a single administrative domain of the TOE*] as to maintain the following conditions: [assignment: *list of conditions on objects*].

Dependencies: No dependencies.

FPR_TRD.3 Allocation of processing activities

Hierarchical to: FPR_TRD.1

FPR_TRD.3.1 The TOE shall be divided in separate, independent, intercommunicating parts (administrative domains) governed by distinct access control and authentication configurations.

FPR_TRD.3.2 The distinct administrative domains of the TOE shall explicitly request access to data stored on other parts of the TOE to be granted access to it.

FPR_TRD.3.3 The TSF shall ensure that [assignment: *list of operations*] shall be performed by different administrative domains of the TOE, so that the following conditions are maintained: [assignment: *list of conditions on operations*].

Dependencies: No dependencies.