ETSI TS 103 485 V1.1.1 (2020-08)

TECHNICAL SPECIFICATION

CYBER;
Mechanisms for privacy assurance and verification

Reference

DTS/CYBER-0013

Keywords

assurance, confidentiality, identification, privacy

ETSI

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - NAF 742 C
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° 7803/88

Important notice

Copyright Notification

Contents

Intellectual Property Rights

Essential patents

IPRs essential or potentially essential to normative deliverables may have been declared to ETSI. The information pertaining to these essential IPRs, if any, is publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: *"Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards"*, which is available from the ETSI Secretariat. Latest updates are available on the ETSI Web server (https://ipr.etsi.org/).

Pursuant to the ETSI IPR Policy, no investigation, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Trademarks

The present document may include trademarks and/or tradenames which are asserted and/or registered by their owners. ETSI claims no ownership of these except for any which are indicated as being the property of ETSI, and conveys no right to use or reproduce any trademark and/or tradename. Mention of those trademarks in the present document does not constitute an endorsement by ETSI of products, services or organizations associated with those trademarks.

Foreword

This Technical Specification (TS) has been produced by ETSI Technical Committee Cyber Security (CYBER).

Modal verbs terminology

In the present document "**shall**", "**shall not**", "**should**", "**should not**", "**may**", "**need not**", "**will**", "**will not**", "**can**" and "**cannot**" are to be interpreted as described in clause 3.2 of the ETSI Drafting Rules (Verbal forms for the expression of provisions).

"**must**" and "**must not**" are **NOT** allowed in ETSI deliverables except when used in direct citation.

1 Scope

The present document defines the means to enable assurance of privacy, using the conventional CIA (Confidentiality, Integrity, Availability) paradigm and with reference to the functional capabilities for privacy protection described in Common Criteria for Information Technology Security Evaluation [1]. The present document addresses privacy assurance within the context of Identity Management following the model described in ETSI TS 103 486 [i.17]. The present document addresses the cases where both transient and persistent identifiers are used, and where identifiers are used in isolation and where identifiers are used in combination.

The mechanisms defined in the present document have been informed by the requirements found in articles and recitals of the General Data Protection Regulation (EU) 2016/679 [i.6] (GDPR) and can be considered in assisting in achieving compliance to the requirements in GDPR.

NOTE: The GDPR contains a very large number of requirements and the present document addresses only a very small number of the technical ones thus the present document is not a solution to the GDPR but where it may assist in addressing specific requirements this has been identified in the body of the present document.

The present document identifies assurance protection levels for privacy and mechanisms for achieving those protection levels. It does not, however, provide any guarantee that application of the mechanisms will prevent abuse of private information and user privacy.

2 References

2.1 Normative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

Referenced documents which are not found to be publicly available in the expected location might be found at https://docbox.etsi.org/Reference/.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long term validity.

The following referenced documents are necessary for the application of the present document.

[1] "Common Criteria for Information Technology Security Evaluation; Part 2: Security functional components"; April 2017 Version 3.1 Revision 5.

NOTE 1: Available from https://www.commoncriteriaportal.org/files/ccfiles/CCPART2V3.1R5.pdf.

NOTE 2: The base text from this reference is also available as ISO/IEC 15408-2.

2.2 Informative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long term validity.

The following referenced documents are not necessary for the application of the present document but they assist the user with regard to a particular subject area.

[i.1] ETSI EG 203 310: "CYBER; Quantum Computing Impact on security of ICT Systems; Recommendations on Business Continuity and Algorithm Selection".

[i.2] ETSI GR QSC 004: "Quantum-Safe Cryptography; Quantum-Safe threat assessment".

[i.3] ETSI TS 102 165-1: "CYBER; Methods and protocols; Part 1: Method and pro forma for Threat, Vulnerability, Risk Analysis (TVRA)".

[i.4] ISO/IEC 29134: "Information technology - Security techniques - Guidelines for privacy impact assessment".

[i.5] INCITS 359-2012: "Information technology - Role Based Access Control".

[i.6] Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation).

[i.7] The OECD Privacy Framework (2013): "Organisation for economic co-operation and development".

NOTE: Available from http://www.oecd.org/sti/ieconomy/oecd_privacy_framework.pdf.

[i.8] ETSI TS 102 165-2: "CYBER; Methods and protocols; Part 2: Protocol Framework Definition; Security Counter Measures".

[i.9] ETSI TS 103 532: "CYBER; Attribute Based Encryption for Attribute Based Access Control".

[i.10] NIST Special Publication 800-53 (Rev. 4): "Security Controls and Assessment Procedures for Federal Information Systems and Organizations".

[i.11] Article 29 Data Protection Working Party: "Opinion 05/2014 on Anonymisation Techniques", Adopted on 10 April 2014.

NOTE: Available from https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2014/wp216_en.pdf.

[i.12] ETSI EN 300 392-7: "Terrestrial Trunked Radio (TETRA); Voice plus Data (V+D); Part 7: Security".

[i.13] ETSI TS 123 401: "LTE; General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access (3GPP TS 23.401)".

[i.14] ANSI X9.42: "Public Key Cryptography For The Financial Services Industry: Agreement Of Symmetric Keys Using Discrete Logarithm Cryptography".

[i.15] IETF RFC 2631: "Diffie-Hellman Key Agreement Method".

[i.16] ISO/IEC 11770-3: "Information technology - Security techniques - Key management - Part 3: Mechanisms using asymmetric techniques".

[i.17] ETSI TS 103 486: "CYBER; Identity Management and Discovery for IoT".

[i.18] ETSI TR 103 370: "Practical introductory guide to Technical Standards for Privacy".

[i.19] ETSI TR 187 010: "Telecommunications and Internet converged Services and Protocols for Advanced Networking (TISPAN); NGN Security; Report on issues related to security in identity imanagement and their resolution in the NGN".

3 Definition of terms, symbols and abbreviations

3.1 Terms

For the purposes of the present document, the following terms apply:

personal data: data relating to the natural person as per the GDPR

private data: data that an entity wants to restrict distribution of

3.2 Symbols

Void.

3.3 Abbreviations

For the purposes of the present document, the following abbreviations apply:

ABAC	Attribute Based Access Control
ABE	Attribute Based Encryption
CBC	Cipher Block Chaining
CIA	Confidentiality Integrity Availability
CP-ABE	Ciphertext Policy - Attribute Based Encryption
CTR	Counter
DAC	Discretionary Access Control
DLT	Distributed Ledger Technology
DPIA	Data Protection Impact Assessment
ECB	Electronic Code Book
GDPR	General Data Protection Regulation
ICT	Information Communications Technology
INCITS	InterNational Committee for Information Technology Standards
IP	Internet Protocol
KP-ABE	Key Policy - Attribute Based Encryption
MAC	Mandatory Access Control
OECD	Organization for Economic and Cultural Development
PII	Personally Identifying Information
PKC	Public Key Cryptography
PKG	Private Key Generator
PKI	Public Key Infrastructure
RBAC	Role Based Access Control
ToE	Target of Evaluation
TSF	ToE Security Function

4 Core concepts of privacy and privacy protection

4.1 Introduction

The guidance given in ETSI TR 103 370 [i.18] is formalized in the present document for application in ICT systems. This aims to address the requirement from GDPR to protect those aspects of a natural person's rights and freedoms which encompass privacy and Personally Identifying Information (PII).

The General Data Protection Regulation (EU) 2016/679 [i.6] (GDPR) was adopted on 27 April 2016 and entered into force on 25 May 2018. Whilst the GDPR has a focus on data protection where data relates to a natural person it is not privacy per-se that GDPR seeks to protect. The GDPR protects aspects of a natural person's rights and freedoms which encompass privacy and PII. Thus, whilst data protection in general is a major objective for the ICT domain, the GDPR requires that any data that can be related, or relatable, to a natural person is protected.

NOTE 1: The preceding directive 95/46/EC to the GDPR was repealed on 25 May 2018 and any references made anywhere to 95/46/EC are automatically assumed to refer to the GDPR [i.6] after this date.

Dictionary definitions of "privacy" are not sufficiently nuanced to capture the emotional content of what privacy means to an individual.

EXAMPLE 1: One common definition of privacy is "a state in which one is not observed or disturbed by other people".

EXAMPLE 2: Synonyms for privacy include the following list of terms: seclusion, privateness, solitude, isolation, retirement, peace, peace and quiet, peacefulness, quietness, lack of disturbance, lack of interruption, freedom from interference, sequestration, reclusion.

Privacy in ICT networks is not well defined either, therefore the present document summarizes the general concept of privacy as it relates to the ICT context using a concept relationship diagram (Figure 1). Private data can be understood as anything an entity wants to restrict distribution of, while personal data relates to the natural person as per the GDPR. Certain forms of identifying data can be directly controlled to restrict availability, whereas some, such as observed behaviour may be less likely to be controlled but the ability to link observations to a specific person should be protected in such a way that the link is weak. Similarly, observation of use of a device can be a vector for identification of an individual and the present document considers means to weaken such links.

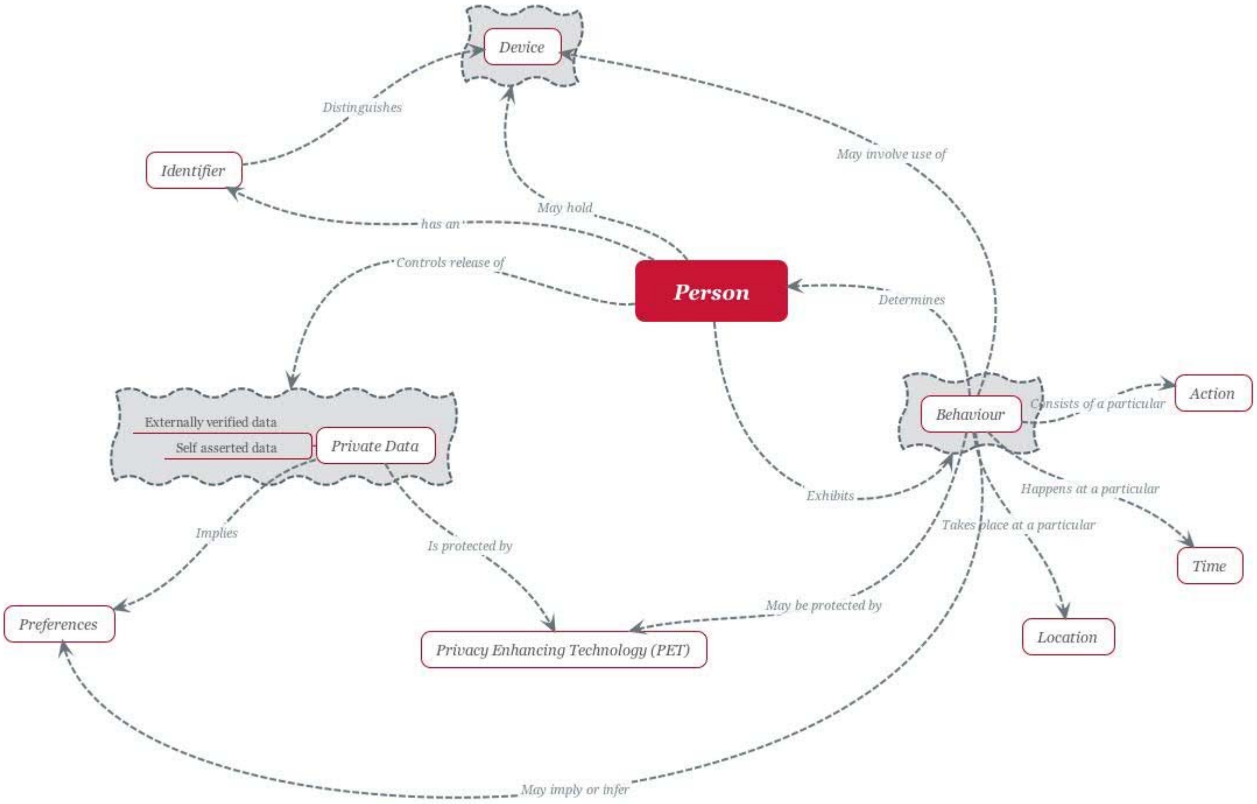

Figure 1: Concept relationship diagram for data relationships that can impact privacy

It is asserted that ICT networks both store and transfer data in a distributed manner across different entities and geographies. This data can be personal, or be linked in such a way as to be considered personal, or private. ICT systems need to exchange data to allow them to operate. Such systems can be required to account for that operation where the data needed can, of necessity, be linked to an individual or closely associated to an individual; thus, ICT systems come under the legal requirements of the GDPR where applicable as per Article 3 of the GDPR [i.6].

NOTE 2: Sometimes data related to a person is private but can also be in the public domain. Even in this case there is a convention often enshrined in law, that those viewing such data will not use it against the person or use it in a new context without reference to the person. This is a social obligation and goes beyond the specific conventions of the GDPR [i.6].

The central entity in Figure 1 is the Person, which the GDPR [i.6] terms the data subject. Within the GDPR the scope of data protection is identified by the relationship between the data subject and each of the data controller (defined in GDPR as controller) and the data processor (defined in GDPR as processor). This is shown in Figure 2 (also given in ETSI TR 103 370 [i.18]) which introduces the concept of explicit consent to the use of private or personal data. It is however clear that explicit consent is only one of the basis for legal processing of personal data cited in GDPR [i.6], Article 6, all of which require policy statements and policy compliance by the data controller and data processor entities.

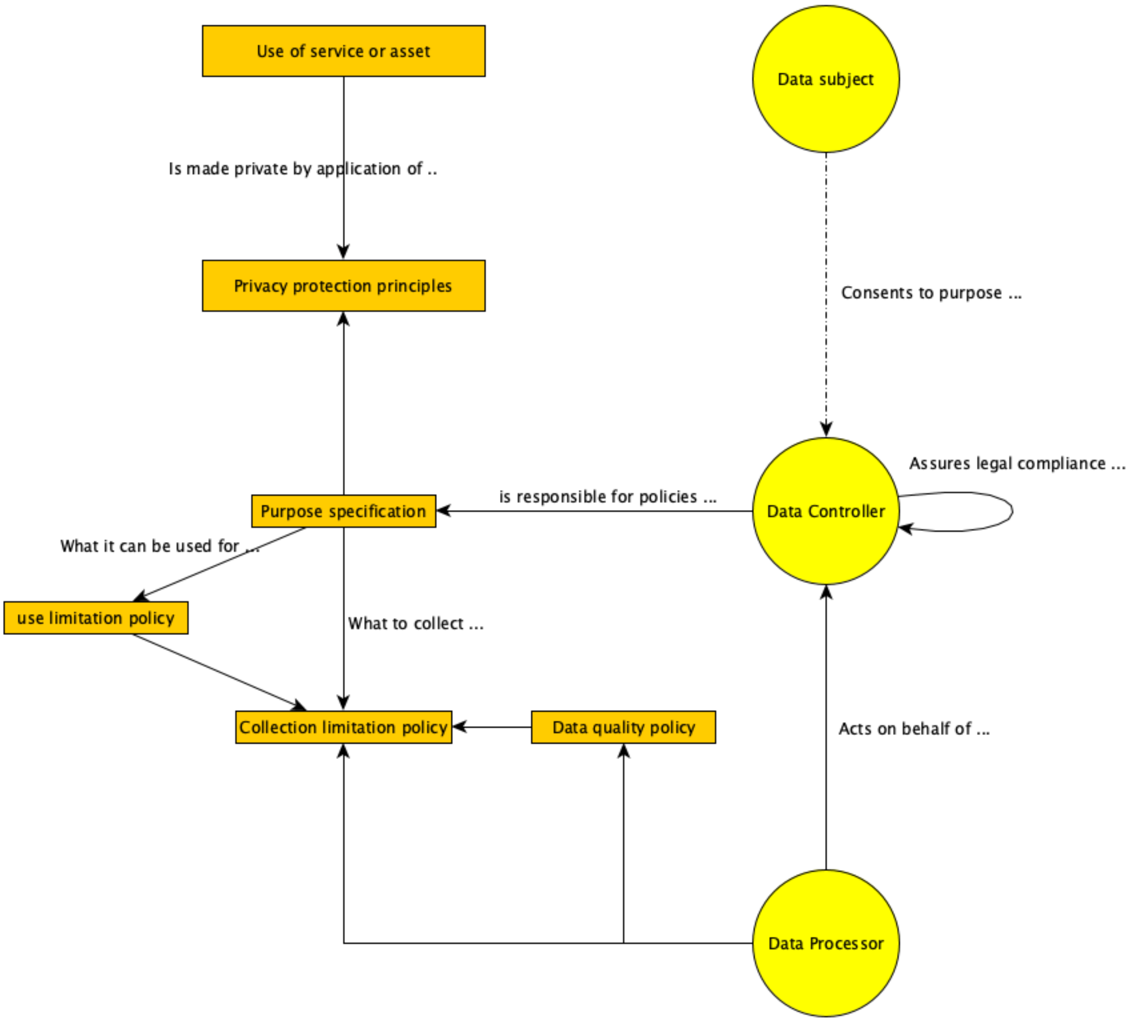

NOTE 1: Data required to enable a service can be greater than that consumed by the service itself but can include data to enforce obligations arising from use of the service, e.g. for payment.
NOTE 2: The data controller has responsibility for assurance of the lawful basis of any processing in addition to the subset of cases that require explicit consent from the data subject.
NOTE 3: The data controller is responsible for establishing the policies that implement each of the core data protection principles, and the data processor is responsible for following such controller-designed policies as a third party to the controller. The processor's responsibilities should be established by contract, as in standard or model contractual clauses when such processor is not located in a geographic area that observes the GDPR or has established adequacy under the European Commission's requirements.
NOTE 4: Consent from the data subject to allow processing is not the only form of lawful permission to allow processing but is the only one requiring direct input from the data subject hence in the figure the relationship is shown by a broken line.

Figure 2: GDPR actors mapped to core data protection principles outlined in clause 4.3

The detail of the data protection principles is examined in each of ETSI TR 187 010 [i.19] and ETSI TR 103 370 [i.18] and examined for the present document in clause 4.2.

Unfortunately, Figure 2 is over-simplistic, it implies an atomic relationship of data subject and data controller, i.e. the relevant lawful basis for processing is sufficient to address all possible uses of data within the accessed service. In many deployments the service that is offered, particularly ICT services, implies access to many sub-services that are not necessarily directly associated to the purpose specified by the data controller.

4.2 Privacy Impact Assessment (PIA)

Data Protection Impact Assessments (DPIAs) help to evaluate, the origin, nature, particularity and severity of risk presented to the rights and freedoms of natural persons when processing operations are implemented. The outcome of the assessment should be used to determine the appropriate measures to be implemented. The risk analysis methods given in ETSI TS 102 165-1 [i.3] and ISO/IEC 29134 [i.4] can be used to perform the DPIA.

4.3 Core data protection principles

The core protection principles for use of private data are already extant in the OECD Guidelines [i.7] and are formalized in each of ETSI TR 187 010 [i.19] and ETSI TR 103 370 [i.18]:

- Collection limitation principle (see also GDPR [i.6], Articles 6, 5.1.a)
- Data quality principle (see also GDPR [i.6], Articles 5.1.d, 5.1.e)
- Purpose specification principle (see also GDPR [i.6], Article 5.1.b)
- Use limitation principle (see also GDPR [i.6], Article 5.1.b)
- Security safeguards principle
- Openness principle
- Individual participation principle
- Accountability principle (see also GDPR [i.6], Article 5.2)
- Equality of regime principle

In addition to the explicitly cited mapping of GDPR articles to the principles above the general recitals and articles of GDPR [i.6] act in support of these general principles.

4.4 Conventional CIA protections

The Confidentiality Integrity Availability (CIA) paradigm provides for mechanisms to assure each of the CIA attributes. Frameworks and patterns for techniques that apply the CIA approach can be found in ETSI TS 102 165-2 [i.8] and are adopted in the remainder of the present document.

5 Privacy protection measures from Common Criteria

5.1 Introduction and summary of measures in Common Criteria

NOTE: This clause contains text from Common Criteria Part 2 [1] to give clarity in the processing of data to provide privacy protection. The texts from these referenced documents are intended for re-use and thus have been treated as templates and modified from the source for the application to the present document.

The Common Criteria functional capabilities defined in Common Criteria Part 2 [1], address privacy protection in the class FPR, Privacy, that aims to provide a user with protection against discovery and misuse of identity by other users. The capabilities in the FPR class need to be considered alongside the capabilities in the Identification and Authentication class (FIA) and in the audit class (FAU) of Common Criteria Part 2 [1]. The measures are outlined below and the detail considerations for assignment with regards to what aspects of private or Personally Identifiable Information (PII) are protected, and how the metrics for the resulting protection are assured.

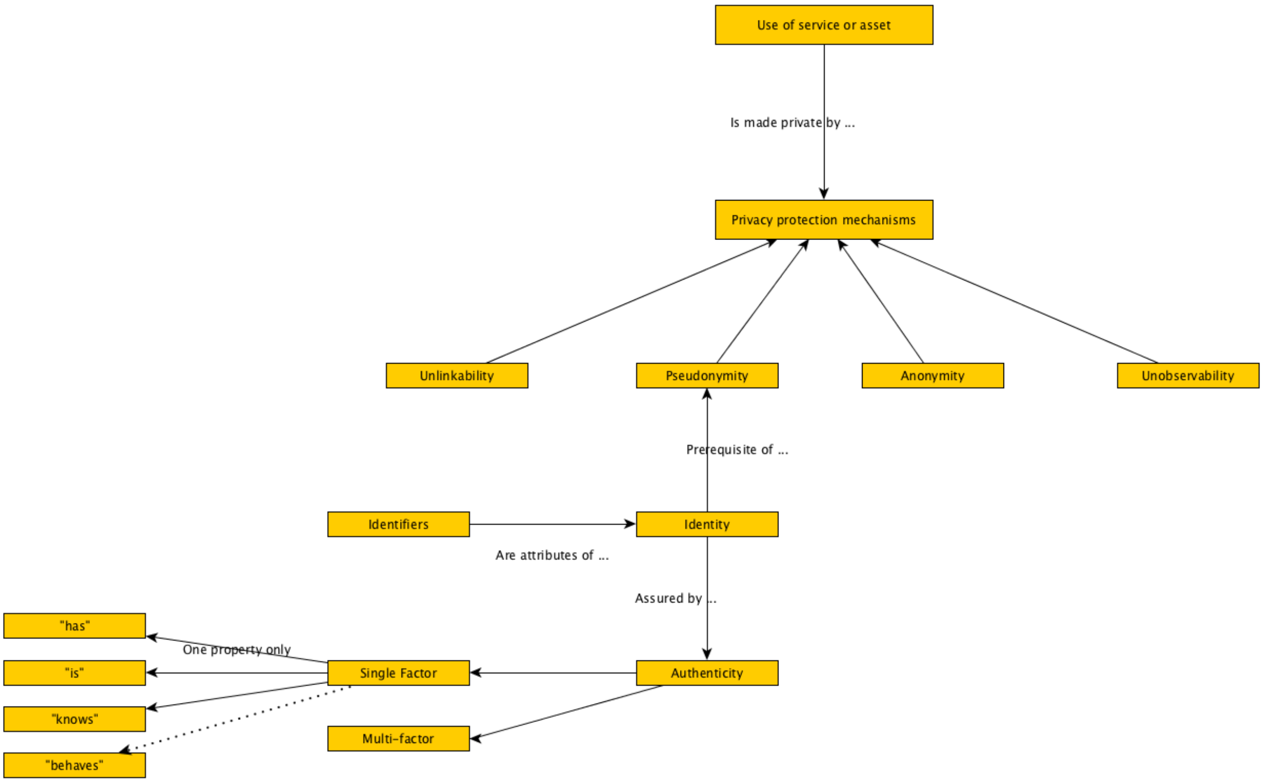

Figure 3: Simplified view of relations between techniques that can protect privacy, taken from Common Criteria

Figure 3 takes the view that the use of a service or access to an asset is made private by specific privacy protection mechanisms. The model that lies behind Figure 3 is that the user is offered various levels of protection by the way in which the service or asset is designed. Thus if the use of the service or asset requires accountability the designer can choose to implement a pseudonymity service, and in order to do so will implement a formal identity management process that allows the assignment of the pseudonym to a known and authenticated user. In contrast if the service has no requirement for accountability the designer can choose to make the use of the service anonymous, requiring no capture of identifying data from the user and thus no framework for identity management and authentication.

EXAMPLE 1: A service requiring accountability is any service where a user needs to be billed for the service in which case a pseudonymity capability is provided.

EXAMPLE 2: A web service with no personalized content is an example where the designer provides an anonymity capability for access to the service.

5.2 Interaction with core data protection principles

The data collection, use and deletion policies shall be declared as shown in Figure 2 before applying any specific privacy protection mechanisms.

NOTE: Assured protection cannot be afforded if there is no bounded policy in which data is processed.

5.3 Anonymity

NOTE 1: This clause contains text from Common Criteria Part 2 [1] to give clarity in the processing of data to provide privacy protection. The texts from these referenced documents are intended for re-use and thus have been treated as templates and modified from the source for the application to the present document. Any text that is taken from the reference documents is clearly indicated by being formatted in *italics*.

In Common Criteria [1] the Anonymity class of functional capabilities gives assurance that a user can use a resource or service without disclosure of the user's identity. The Common Criteria defines terms that are re-used in the Tables 1 through 4. The core of these relate to a Target of Evaluation, defined as a bounded environment to which the security claims and analysis apply, and within that are ToE Security Functions (TSFs) that provide the means to give assurance. In the context of GDPR [i.6], Recital 26 data that is anonymous, or strictly anonymized following application of the methods outlined in the present document, is not directly impacted by GDPR. The functional capabilities defined in the Common Criteria Anonymity class may be used (optional), but if used shall be used as indicated in Table 1.

NOTE 2: In Tables 1 through 4 the convention of referring to legitimate parties as Alice and Bob, and referring to an adversary as Eve, is followed.

A distinction is drawn in Table 1 between an individual interacting with a service anonymously and the wider concept of data anonymisation. The primary target of the Anonymity class is the former, i.e. allowing a user to interact with a service anonymously. The slightly wider concept of anonymisation of data where datasets may be used as a tool in re-identification primarily refers to use of direct and indirect data in order to provide an association between a real entity, Alice, and the anonymous identifier.

Table 1: Anonymity class applied to privacy assurance and verification

Shortname	Definition (text in italics comes from Common Criteria part 2 [1])	Comments with respect to mechanisms for privacy assurance and verification
FPR_ANO.1.1 Anonymity	*The TSF shall ensure that* any instance of an adversary, Eve, is *unable to determine the real user name bound* to any instance of Alice.	Anonymity provisions are designed such that actions of Alice are not linkable to an instance of Alice. Indirect means of identification shall be protected in such a way that they are not able to be used as channels to re-identify Alice, e.g. behavioural analysis as shown in clause 4.1, Figure 1.
FPR_ANO.2.1 Anonymity without soliciting information	*The TSF shall ensure that* any instance of an adversary, Eve, *is unable to determine*, by any direct or indirect means, *the real user name bound to* any instance of Alice.	
FPR_ANO.2.2	*The TSF shall provide [**assignment:** list of services] to* any instance of an anonymized Alice *without soliciting any reference to the real user name*.	Providers of services should seek to minimize the set of services offered without identification of the user. Providers are obligated in all cases to comply with appropriate regulations.

When Alice is represented by an anonymous identifier the constraints in Table 1 shall apply. The wider concern of re-identification of Alice shall not be possible, i.e. for Eve to re-identify Alice, by correlating multiple sets of anonymized data. After implementation of any anonymisation countermeasure intended to achieve the functional requirements in Table 1 the following attributes of the resultant data should be present to prevent re-identification of Alice:

- Immunity to Singling out: It shall be infeasible for Eve to isolate some or all records which identify Alice in the dataset.

- Immunity to Linkability: It shall be infeasible for Eve to link *n* records (where *n* is greater than or equal to two (2)) concerning the same data subject (Alice) or a group of data subjects (either in the same database or in two different databases).

- Immunity to Inference: It shall be infeasible for Eve to deduce, with significant probability, the value of an attribute associated to Alice from the values of a set of other attributes.

The processing of multiple anonymized datasets to achieve re-identification is shown in Figure 4. Techniques to mitigate re-identification including the following should be taken into consideration:

- Randomization, wherein data is modified to weaken the link between Alice and her data.
- Generalization, wherein data is structured into wider groupings.

EXAMPLE 1: If the original dataset contains records by specific house number and street name per town this data is generalized to town name.

EXAMPLE 2: If the original dataset contains records of Alice's age the data is generalized to age ranges, i.e. original record stores age as (say) 27 years old, the generalized record stores age as being in the range 20 < age < 30.

More specific approaches to achieve each of randomization and generalization that should be considered include the following:

- k-anonymity, applied within any single dataset should maximize the value of k (i.e. if a database table has n-attributes and some subset of the attributes are anonymized then any search across the remaining attributes should always return at least k records).
- Noise injection, wherein data may be made noisy (i.e. modifying the anonymized data set to provide uncertainty on any particular data value) or additional data records are added to mask real data.
- Differential privacy, wherein an anonymized and noisy dataset is further subdivided by forms of grouping (somewhat similar to using generalization on groups of data) of data. An algorithm or search on data is considered as differentially private if Eve seeing its output (the search result) cannot tell if Alice's information was used in the computation (or was included in the searched source).

NOTE 3: The EU Data Commissioners Article 29 advisory group has drafted an opinion on techniques to achieve anonymisation. Designers are invited to consult this opinion prior to implementing any of the techniques in this list [i.11].

5.4 Pseudonymity

NOTE: This clause contains text from Common Criteria Part 2 [11] to give clarity in the processing of data to provide privacy protection. The texts from these referenced documents are intended for re-use and thus have been treated as templates and modified from the source for the application to the present document. Any text that is taken from the reference documents is clearly indicated by being formatted in *italics*.

In Common Criteria [1] the Pseudonymity class of functional capabilities gives assurance that a user can use a resource or service without disclosing their user identity, but can still be accountable for that use (i.e. the user can be reidentified if necessary). The role of pseudonymity as a risk mitigation method for protection of personal data is cited in GDPR [i.6], Recitals 26, 28, 29, 75, 85, 175 and in Articles 6.4.e, 25.1, 32.1.a, 40.2.d, 89.1.

Table 2: Pseudonymity class applied to privacy assurance and verification

Shortname	Definition (text in italics comes from Common Criteria part 2 [1])	Comments with respect to mechanisms for privacy assurance and verification
FPR_PSE.1.1	*The TSF shall ensure that* any instance of an adversary, Eve, is *unable to determine the real user name bound to* the pseudonymous representation of any real user Alice.	
FPR_PSE.1.2	*The TSF shall be able to provide* at least one (1) *alias of the real user name* (Alice) *to* each service used by Alice.	
FPR_PSE.1.3	*The TSF shall [**selection**, choose one of: determine an alias for a user, accept the alias from the user] and verify that it conforms to the [**assignment:** alias metric]*.	
FPR_PSE.2.1	*The TSF shall ensure that* any instance of an adversary, Eve, is *unable to determine the real user name bound to* the pseudonymous representation of any real user Alice.	Whilst the wording of the requirements in FPR_PSE.2.x are identical to those of FPR_PSE.1.x, PSE.2.x refers to "reversible pseudonymity" which requires the system to allow for reverse lookup to happen (in FPR_PSE.2.4)
FPR_PSE.2.2	*The TSF shall be able to provide* at least one (1) *alias of the real user name* (Alice) *to* each service used by Alice	
FPR_PSE.2.3	*The TSF shall [**selection**, choose one of: determine an alias for a user, accept the alias from the user] and verify that it conforms to the [**assignment:** alias metric]*.	
FPR_PSE.2.4	Pseudonyms shall only be reversible by authorized parties.	
FPR_PSE.3.1	*The TSF shall ensure that* any instance of an adversary, Eve, is *unable to determine the real user name bound to* the pseudonymous representation of any real user Alice.	As for FPR_PSE.1.x
FPR_PSE.3.2	*The TSF shall be able to provide* at least one (1) *alias of the real user name* (Alice) *to* each service used by Alice.	
FPR_PSE.3.3	*The TSF shall [**selection**, choose one of: determine an alias for a user, accept the alias from the user] and verify that it conforms to the [**assignment:** alias metric].*	
FPR_PSE.3.4	*The TSF shall provide an alias* to Alice which shall not be linkable to any alias previously given to Alice other than in explicitly identified conditions.	

EXAMPLE: Pseudonymity is used in the Temporary Mobile Subscriber Identity (TMSI) in 2G and 3G cellular telephony (see ETSI TS 123 401 [i.13]), the Alias Short Subscriber Identity (ASSI) and Encrypted Short Identity (ESI) mechanisms in TETRA (see ETSI EN 300 392-7 [i.12]), and provisioned by aliases in online systems including those used in opinion forums, and gaming sites.

5.5 Unlinkability

NOTE: This clause contains text from Common Criteria Part 2 [1] to give clarity in the processing of data to provide privacy protection. The texts from these referenced documents are intended for re-use and thus have been treated as templates and modified from the source for the application to the present document. Any text that is taken from the reference documents is clearly indicated by being formatted in *italics*.

In Common Criteria [1] the Unlinkability class of functional capabilities give assurance that a user can make multiple uses of resources or services without adversaries being able to link these uses together.

Table 3: Unlinkability class applied to privacy assurance and verification

Shortname	Definition (text in italics comes from Common Criteria part 2 [1])	Comments with respect to mechanisms for privacy assurance and verification
FPR_UNL.1.1	*The TSF shall ensure that* any instance of an adversary, Eve, is *unable to determine whether* observable actions in the ToE were caused by the same instance of Alice	

EXAMPLE: Unlinkability using the ESI feature in TETRA [i.12] inhibits the ability of the adversary, Eve, to use traffic analysis to determine the activity of any individual user.

5.6 Unobservability

NOTE: This clause contains text from Common Criteria Part 2 [1] to give clarity in the processing of data to provide privacy protection. The texts from these referenced documents are intended for re-use and thus have been treated as templates and modified from the source for the application to the present document. Any text that is taken from the reference documents is clearly indicated by being formatted in *italics*.

In Common Criteria the Unobservability class of functional capabilities give assurance that a user can make use of resources or services without others, especially third parties, being able to observe that the resource or service is being used. The concept of unobservability can contradict application of the accountability and openness principles outlined in clause 4.3. In some cases such as the application of certain management roles, and the specific instances of lawful interception, any provision of an unobservability measure may be overridden by authorized users (this is specifically applicable to FPR_UNO.3.1 and FPR_UNO.4.1 in Table 4).

Table 4: Unobservability class applied to privacy assurance and verification

Shortname	Definition (text in italics comes from Common Criteria part 2 [1])	Comments with respect to mechanisms for privacy assurance and verification
FPR_UNO.1.1	*The TSF shall ensure that* any instance of an adversary, Eve, *is unable to observe the operation [**assignment:** list of operations] on [**assignment:** list of objects] by* any representation of any real user Alice.	As noted in the text introducing this table the purpose of Unobservability is to ensure that it is not possible to identify if resources are in use by Alice.
FPR_UNO.2.1	*The TSF shall ensure that any instance* of an adversary, Eve, *is unable to observe the operation [**assignment:** list of operations] on [**assignment:** list of objects]* by any representation of any real user Alice.	
FPR_UNO.2.2	*The TSF shall allocate the [**assignment:** unobservability related information] among different parts of the TOE such that the following conditions hold during the lifetime of the information: [**assignment:** list of conditions].*	
FPR_UNO.3.1	*The TSF shall provide [**assignment:** list of services] to [**assignment:** list of subjects] without soliciting any reference to [**assignment:** privacy related information].*	
FPR_UNO.4.1	*The TSF shall provide [**assignment:** set of authorized users] with the capability to observe the usage of [**assignment:** list of resources and/or services].*	

Provision of an unobservability service is often very difficult to justify in normal commercial networks and services and thus the capabilities described in Table 4 are offered for completeness of review of the functional capabilities described in Common Criteria and should only be applied with caution.

The implementation of unobservability results in prevention of the ability of Eve to identify Alice through traffic analysis and processing load analysis with additional measures to give assurance of confidentiality.

6 Relationship cardinality for privacy assurance

6.1 Overview

The number of entities in a relationship determines to some extent the difficulty of managing (private) data within that relationship. The cardinality of the relationship between Alice and Bob shall be determined in advance of any data exchange. As the number of entities in a relationship grows the risk of exposure of private data relating to any participant in the relationship also grows. Where private data is given to unknown parties the risk of exploit of that data is considered higher than when sharing private data amongst known parties.

In defining a data policy the data controller shall identify the cardinality of relationships of users to data and shall make appropriate selection of cryptographic protections for data as a result. The remainder of this clause identifies specific considerations to be taken by the data policy designer.

The cardinality of relationships to be protected can be used to select the crypto-system architecture. A summary of the characteristics of the relationships with respect to the risk of exposure of private data follows:

- One to one (1:1) relationships:
    - In this relationship there is only one instance of each of Alice and Bob;
    - Alice is assumed to know the identity of Bob and Alice should be able to authenticate the identity of Bob (and vice versa).
- One to many (1:m) and many to one (m:1) relationships:
    - In this case there is one instance of Alice and many instances of Bob (Bob can be a member of a group and cannot be uniquely distinguished), in such cases Alice does not need to know which particular instance of Bob she is communicating with;
    - Alice is known to all instances of Bob and Bob should be able to authenticate the identity of Alice.
- Many to many (m:n) relationships:
    - In this case many instances of Alice are able to communicate with many instances of Bob.

6.2 One to one

Here, there is only one instance of each of Alice and Bob and they can exchange a secret in advance, perhaps out-of-band pre-shared keys or through public-key authentication; the present document does not consider how this shared key has been distributed ahead of time. If a shared key has been previously distributed, Alice and Bob's communications will be protected by conventional symmetric cryptography.

As data is shared only within the relationship, each party knows any decrypted data shared more widely has been shared by themselves or the other party. This is the simplest case of managing of private data in terms of cardinality.

Where Alice and Bob authenticate each other by means of pre-shared keys the key shall not be directly exposed, rather each party shall prove to the other that they have the key. Authentication shall be achieved using challenge-response methods as outlined in ETSI TS 102 165-2 [i.8].

If Alice and Bob do not have a pre-shared secret, they should initiate an authenticated connection using identified asymmetric key exchange and generation methods as outlined in ETSI TS 102 165-2 [i.8]. One such method is the Diffie-Hellman key agreement protocol [i.14], [i.15] and [i.16].

6.3 One to many, many to one

In a one-to-many or many-to-one data sharing relationship, Alice shares data with multiple other parties, but these other parties do not communicate with each other directly. This is similar to the one-to-one case, but with multiple instances of Bob referring to one Alice.

In larger systems an infrastructure may be used to support the distribution of public keys via public key certificates. These certificates include an assertion by a trusted 3rd party that the associated private key is bound to an attribute of the public key holder.

EXAMPLE: A Public Key Infrastructure is used to distribute keys and the associated trust framework.

NOTE: Privacy aspects of public key certificates in a PKI are considered in clause A.2 of the present document.

6.4 Many to many

In this case, a user community can be large, but for some certain class of data, there is a subgroup of the community that needs to be able to share that data securely without others accessing it.

EXAMPLE: Database access or group communications are often sub-divided according to role or privilege level. Other privilege distinctions include the ability to modify the data or have read-only access.

This cardinality is often managed with a central key management authority which can create a group key and distribute it to all members in that group, using a secure channel to each member.

Many to many relationships are composed of data created and shared by multiple instances of an arbitrary set of Alices to multiple instances of Bob, perhaps across different platforms or databases. This situation lends itself to Attribute Based Encryption (ABE) (see clause A.3 for further detail).

A common use case for many to many relationships is where an object (file, executable code, etc.) needs to be (and remains) encrypted at all times. Alice is the initiator of the file, Bob and Charles can modify it whilst it remains encrypted, but without Bob or Charles being able to extract the plain-text of the object whilst assuring that legitimate parties can perform authorized actions. This use case is mathematically complex to assure. Solutions do exist: Homomorphic encryption; and, Attribute based encryption, can satisfy this use case. Difficulties remain however if the set of legitimate parties are unknown when setting out the access policy. This is an area of active cryptographic research.

7 Access control mechanisms in protecting private data

7.1 General provisions

The use of access control shall be used to counter threats and their associated threat agents that lead to attacks of unauthorized use, disclosure, modification, destruction and denial of service.

NOTE: Access control counter-measures act independently of authentication and confidentiality countermeasures but are used to determine if Alice (ideally confirmed through authentication) is authorized to use the service or access the data requested. FIPS SP 800-53 [i.10] describes many modes of access control.

Alice shall act as the party seeking access to the data, Bob shall act as the party granting authorization. In all cases Bob shall establish the access control policy including determination of any restrictions placed on Alice once Alice has access to the data. Bob shall explicitly declare if Alice has discretion to choose how to use data (see clause 7.2), or if Alice is bound to policies set by Bob (see clause 7.3).

Clauses 7.2 to 7.5 provide information on existing access control models.

If designing with a view to assurance, the designer should take note of the access control functional capabilities defined in the User Data Protection class in clause 11 of Common Criteria [1].

The data access model shall be defined in the data access policy established by the data owner (Bob in this example). The access control policy shall address the following aspects:

- the subjects under control of the policy (i.e. the set of Alices that Bob defines and manages access for);
- the objects under control of the policy (i.e. the set of things that Bob is controlling access to); and
- the operations among controlled subjects and controlled objects that are covered by the policy.

7.2 Discretionary Access Control (DAC)

In the Discretionary Access Control (DAC) model the entity responsible for the asset (e.g. the private data) determines who has access to it. However once Alice has retrieved the data, Alice can use that data at her discretion, in other words if Bob is the owner of the data and grants Alice access then Bob has no role in what Alice does with that data.

NOTE: There can be some constraints set by Bob via contract, or by regulation, that impose a fair use obligation on, or directly restrict use by, Alice that is not directly enforceable by Bob thus maintaining a DAC model.

7.3 Mandatory Access Control (MAC)

Mandatory Access Control (MAC) extends the DAC model by enforcing restrictions on actions Alice can take on data she has been granted access to.

Technical enforcement of MAC is more difficult than for DAC and can require that Alice and Bob share mutually compatible data access and usage policies. Details of how to achieve MAC enforcement (by Bob of Alice's use of data after accessing that data) are not addressed by the present document.

7.4 Role Based Access Control (RBAC)

NOTE: The models of DAC and MAC apply to RBAC, where RBAC only determines the means by which Alice is offered access.

In RBAC access to the asset is granted only to entities holding a particular role.

EXAMPLE 1: Finance records can be restricted to someone with the role "finance officer" such that an access rule can be written:

```
If role = "finance officer" then grant-access, else deny-access
```

A role can be considered as PII. How easily a role can be used to identify a single individual (unique identification) depends on how many individuals are able to take the role.

EXAMPLE 2: The role of "Prime Minister of the UK" refers to only one person at any one time, so the role is obviously PII.

EXAMPLE 3: The role of "civil servant of the UK" refers to a grouping of approximately 400 000 people, so the role is not obviously PII.

A framework for RBAC is described in INCITS 359-2012 [i.5].

7.5 Attribute Based Access Control (ABAC)

NOTE 1: The models of DAC and MAC apply to ABAC, where ABAC only determines the means by which Alice is offered access.

ABAC is the generalization of RBAC where rather than just using the attribute of role any (arbitrary) set of attributes are used to control access. The greater the set of attributes used in ABAC then the greater the likelihood of identifying an individual, hence the attribute set of ABAC can be PII.

NOTE 2: Whilst Attribute Based Encryption as defined in ETSI TS 103 532 [i.9] offers access control for encrypted data it is a special case of confidentiality protection and is not addressed here for access control.

ABAC uses binary equations to establish access. Sets of policies and rules can be combined in any reasonable Boolean expression as shown in the Figures 4 and 5.

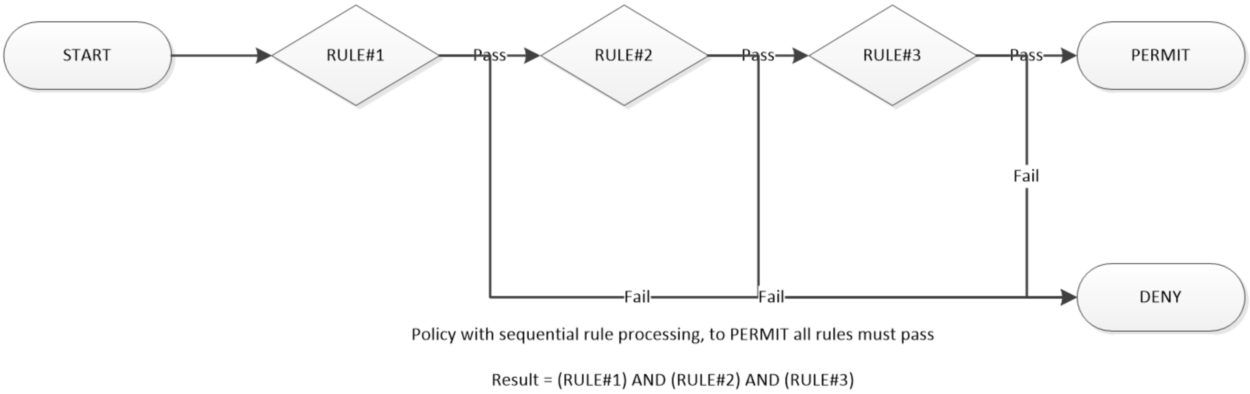
START
RULE#1
Pass
RULE#2
Pass
RULE#3
Pass
PERMIT
Fail
Fail
Fail
DENY
Policy with sequential rule processing, to PERMIT all rules must pass
Result = (RULE#1) AND (RULE#2) AND (RULE#3)

Figure 4: Access control combining rules where all rules have to pass to grant access

A rule can be based on Role (RBAC) or the values of attributes (ABAC).

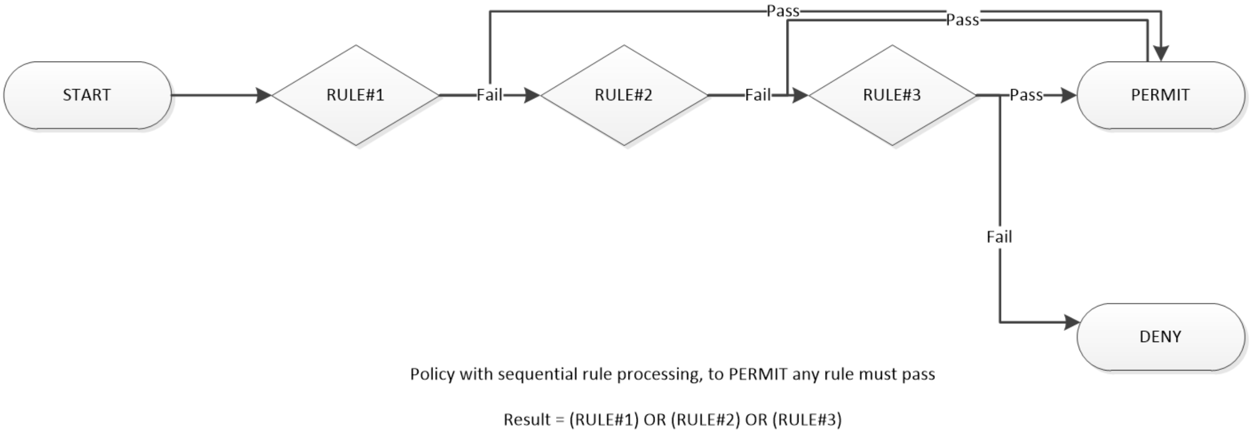
Pass
Pass
START
RULE#1
Fail
RULE#2
Fail
RULE#3
Pass
PERMIT
Fail
DENY
Policy with sequential rule processing, to PERMIT any rule must pass
Result = (RULE#1) OR (RULE#2) OR (RULE#3)

Figure 5: Access control combining rules where any rule has to pass to grant access

As indicated above, in like manner to RBAC there is a risk that the combination of attributes used in the ABAC ruleset may be considered as PII. The use of ABAC is often intended to add additional constraints in access control, such as limiting the time at which access is granted, or where access is granted from (geo-location or virtual (network) location). The example that follows uses the rule combining model shown in Figure 4.

EXAMPLE: Access is restricted to a user with role "finance-officer" attempting to access the resource from the internal network during working hours only, thus:

```
If ("finance officer") AND ("working hours") AND ("Internal network")
    then grant-access,
    else deny-access
```

8 Integrity control mechanisms

8.1 Overview

Integrity control mechanisms primarily enable the accountability principle and data quality principle as outlined in clause 4.3 and in each of ETSI TR 187 010 [i.19] and ETSI TR 103 370 [i.18]. The purpose of integrity control is to verify that a document, audit trail or system cannot be modified without such changes being visible.

In addition, integrity control mechanisms (using the natural language definition of the word) are a natural consequence of the GDPR for implementing rights and freedoms other than privacy (in particular right to access (GDPR [i.6], Article 15) and right to restriction of the processing (GDPR [i.6], Article 18)). The role of integrity control in this case is closer to the concept of provenance, i.e. knowing where something came from, as distinct from knowing that data is complete (whole and undivided), and thus the chain of integrity (including that of provenance) has to be analysed for any privacy implication.

Integrity protection functional capabilities are described in Common Criteria [1] for both data in storage and data in transit. Due attention to the use of functional capabilities described in Common Criteria classes FDP_ITT, FDP_SDI and FDP_UIT should be taken. In particular the following controls should be implemented for a layered transmission system in which layer-N relies on Layer-(N-1) to assure secure transfer of data:

- FDP_ITT.3.1: Entities at Layer N-1 shall monitor user data transmitted between physically-separated parts of the system for integrity errors.
- FDP_ITT.3.2: Upon detection of a data integrity error, the receiving entity at Layer N-1 shall discard the affected data and inform the dependent entity at Layer-N that a data integrity error was detected.

Mechanisms that may be used in the detection of integrity errors are considered in clauses 8.2 to 8.5.

8.2 Cryptographic hashes

Any modern cryptographic hash algorithm produces a fingerprint of a file with the core property that it is infeasible, given only the hash, to retrieve the file that the hash is a fingerprint of, or to create a file that produces the same hash. Similarly it is infeasible to modify the original file in such a way that the hash is itself unchanged (i.e. any modification to the file will be visible through the hash). A final property of hashing functions is that it is infeasible to create two different files that produce the same hash.

The following recommendations apply to any cryptographic hash supporting integrity control:

- One way function recommendation:
  - Given a hash h it should be difficult to find any message m such that $h = \text{hash}(m)$
- Collision resistance recommendation:
  - Given an input $m1$ it should be difficult to find another input $m2$ such that $m1 \neq m2$ and $\text{hash}(m1) = \text{hash}(m2)$
- Strict Avalanche Criterion:
  - if 1 bit in message m changes all other bits should change with 50% probability
- Bit Independence Criterion:
  - output bits j and k should change independently when any single input bit i is inverted, for all i, j and k

The algorithms used to generate a hash function should be of similar strength to any other cryptographic operation in the system.

EXAMPLE: Hashes of medical records ensure the quality and fullness of the data, so that medical professionals can be sure they are seeing information that has not been manipulated.

8.3 Distributed ledgers

Distributed ledger technology, of which blockchain is one example, provides an accurate audit trail of what data has been accessed and when. There is no central administrator, and a peer-to-peer network is required; for this and other reasons with respect to the privacy protection principles outlined in clause 4, and to wider aspects of GDPR, there are aspects of DLT that do not map: see clause A.5 for more details.

EXAMPLE 1: The right to modify false or misleading data is not supported in DLT as once accepted into the ledger no transaction can be modified.

EXAMPLE 2: The GDPR expects a single authority, the data controller, to be accessible to establish and manage the policies and this appears incompatible with the DLT model.

8.4 Signatures

Cryptographic signatures can contribute to the assurance of integrity protection mechanisms where, for example, two or more parties have communicated before trust establishment is complete, and need assurance that all parties have received the same data.

EXAMPLE: At the end of a key exchange or other trust establishment protocol, each party can hash the concatenated exchanged messages, and sign and send the resulting value, provably binding it to their identity with this signature.

8.5 Privacy and integrity controls

Cryptographic hashing mechanisms are not identity revealing of themselves, however where a hash is exchanged to mask other data (say user name and password) it can be uniquely associated to an individual even if no explicit personal data is visible.

EXAMPLE: If a user-name and password are input to a hash function the hash uniquely identifies the username and password combination (as it "fingerprints" the input).

Signatures are cryptographically bound to an identity, so in many applications are personally identifying.

9 Confidentiality mechanisms

In a simplified view of confidentiality Alice and Bob share data in such a way that any adversary, Eve, cannot access the meaning of the content. The most common approach to achieve confidentiality is keyed encryption; see also Annex A. Niche techniques, such as steganography, seek to conceal the existence of the content.

Of themselves means to achieve confidentiality are not identity revealing, and moreover make the transfer of PII in content impossible to identify. However, the existence of an observable encrypted communication channel can reveal some behavioural aspects of Alice and Bob (i.e. that they wish to keep information they share confidential) but as the use of encryption becomes more-pervasive this threat is reduced.

For protection of PII in content and signalling confidentiality mechanisms, i.e. encryption, should be used. However, the key management mechanisms can provide a side channel to reveal some PII as discussed in clause A.2.

Annex A (informative):
Cryptographic mechanisms for privacy assurance and protection

A.1 Symmetric key cryptography

The nature of symmetric key cryptography is that a single key is used for the cryptographic operations where the key is only known to the corresponding parties. Cryptographic protections afforded using symmetric keys rely upon a trusted point to point relationship (see clause 6).

A.2 Asymmetric key cryptography

The nature of asymmetric key cryptography is that a 2- part key exists in which one part, the secret key, is known only to the holder, and the second part, the public key, can be known to anybody (i.e. made public). The public part can be distributed in a number of ways although for large systems some form of Public Key Infrastructure (PKI) can be required to distribute public keys in the form of Public Key Certificates (PKC). A PKC can reveal additional data over and above the public key itself in order for the consuming party to know when to use the contained public key. In such cases the content of the PKC is addressed in the modelling against the core data protection principles addressed in clause 4 of the present document. A privacy concern that can be overlooked is that if the PKI/PKC design is overlaid on the operational system it can lead to data leakage outside the terms of the original purpose specification. This is a side channel as suggested in clause 9. The latter is a consequence of a PKC making an explicit binding of the public key to an attribute (identity or other) that is asserted and verified by a trusted third party that is not necessarily the data controller of the attached service.

A.3 Attribute keyed cryptography

A.3.1 General overview

Attribute Based Encryption (ABE) is a form of encryption based on a secret sharing scheme mapped to attributes and policies. The decryption of a ciphertext is possible only in combination of a secret key allowing retrieval of the shared secret. The nature of ABE is such that it is for use in closed systems as although the form of ABE is asymmetric (different keys used for each of encryption and decryption) the keys are managed by a single master secret.

If the attributes are identifiers relatable to an individual, then the fact that they are public means that they can be classified as identity revealing (in like manner to conventional asymmetric cryptography outlined in clause A.2). Further, and more detailed, definition of ABE is given ETSI TS 103 532 [i.9].

A.3.2 Ciphertext-Policy ABE

In Ciphertext-Policy Attribute Based Encryption (CP-ABE) a user's secret-key is associated with a set of attributes and a ciphertext specifies an access policy over a defined universe of attributes within the system. Decryption of a ciphertext is only possible if the attributes held by the decrypting user match the policy of the respective ciphertext. Policies can be defined over attributes using conjunctions, disjunctions and (k,n) threshold gates, i.e. k out of n attributes have to be present.

EXAMPLE: For a universe of attributes defined to be {A,B,C,D} and user 1 receives a key to attributes {A,B} and user 2 to attribute {D}. If a ciphertext is encrypted with respect to the policy $(A \wedge C) \vee D$ (the conjunction of A and C, disjunction with D), logically also written as ((A AND C) OR D), then user 2 {D} will be able to decrypt, while user 1 {A,B} will not be able to decrypt.

A consequence of CP-ABE is that it provides implicit authorization, i.e. authorization is included into the encrypted data and only people who satisfy the associated policy can decrypt data.

A.3.3 Key-Policy ABE

KP-ABE is the dual to CP-ABE in the sense that an access policy is encoded into the users secret key, e.g. $(A \wedge C) \vee D$, and a ciphertext is computed with respect to a set of attributes, e.g. {A,B}. In this example the user cannot decrypt the ciphertext but can for instance decrypt a ciphertext with respect to {A,C}.

A consequence of KP-ABE is that it provides delayed authorization, i.e. the authorization is included in the users' secret keys, which can be issued after encryption has taken place.

A.3.4 Collusion resistance

An important property which has to be achieved by both CP- and KP-ABE is called collusion resistance. This property implies preventing distinct users to "pool" their secret keys such that they could together decrypt a ciphertext that none of them could decrypt on their own.

A.4 Identity keyed cryptography

Identity-based cryptography is a form of public-key cryptography in which a publicly known string representing an individual or organization is used as a public key (e.g. an email address, domain name, or a physical IP address). Identity-based systems allow any party to generate a public key from a known identity value (an identifier in the terminology of the present document). A trusted third party, called the Private Key Generator (PKG), generates the corresponding private keys. To operate, the PKG first publishes a master public key, and retains the corresponding master private key (referred to as master key). Given the master public key, any party can compute a public key corresponding to the identity ID by combining the master public key with the identity value. To obtain a corresponding private key, the party authorized to use the identity ID contacts the PKG, which uses the master private key to generate the private key for identity ID.

Where an identifier associated to a known individual or organization is used as the public key the cryptographic mechanism is by design identity revealing.

A.5 Distributed Ledger Technology (DLT)

By default Distributed Ledger Technologies (DLTs) share data with a large number of stakeholders and they co-operatively agree to the secure recording of each transaction and later store sets of transactions in a sealed format where the sealing key is agreed to by all parties, the set of transactions is most often referred to as a block. Subsequent blocks use memory of all previous blocks in calculating the sealing record, hence the derivation for blockchain, somewhat analogous to Cipher Block Chaining modes in encryption. Transactions in distributed ledgers can be stored in clear but protected from manipulation.

The privacy protecting capabilities of DLT are to a large extent dependent on the nature of the transactions that are recorded. The stakeholders in the DLT group, i.e. those who maintain the ledgers, each have knowledge of all transactions in the ledger. Therefore, if a recorded transaction contains PII, then, by default, that transaction and the contained PII is known to all, i.e. it becomes public initially to only the blockchain group, but this can be extended to anyone (see Figure A.1).

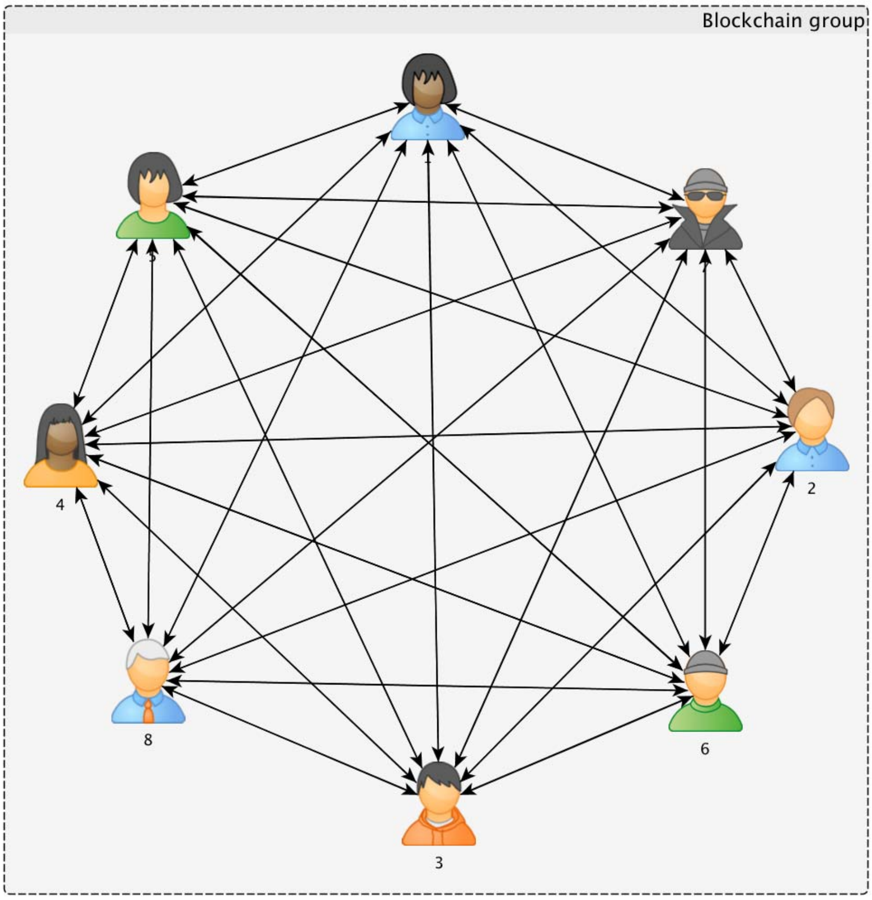

Figure A.1: Scope of knowledge sharing in a distributed ledger or blockchain

The nature of blockchain implies that each of the stakeholders has some trust that the transactions they are asked to record, and to mathematically determine the integrity of, are true records of the claims in the transaction. Thus if in a supply chain a supplier asserts in a ledger record that he shipped 10 widgets from Arkansas to Zimbabwe then it is assumed that this is a true record of what actually happened. If the assertion is false (by malicious or accidental action) then this is not visible in the DLT records, rather the falsehood is maintained across the record.

With respect to the privacy protection principles outlined in clause 4.3 and to wider aspects of GDPR there are aspects of DLT that do not map. In particular the right to modify false or misleading data is not trivial to achieve (this is addressed in the individual participation principle). Furthermore the "use limitation" principle is broken if the content of the ledger is made fully public. The GDPR expects a single authority, the data controller, to be accessible to establish and manage the policies and this appears incompatible with the DLT model.

A.6 Protection against attack modes of cryptography

A.6.1 Encryption modes

Encryption modes that support a Time Variant Parameter between blocks assist in masking the key over time, thus modes such as Electronic Codebook Mode (ECB) ought to be avoided, rather modes such as Cipher Block Chaining (CBC) and Counter (CTR) are preferred.

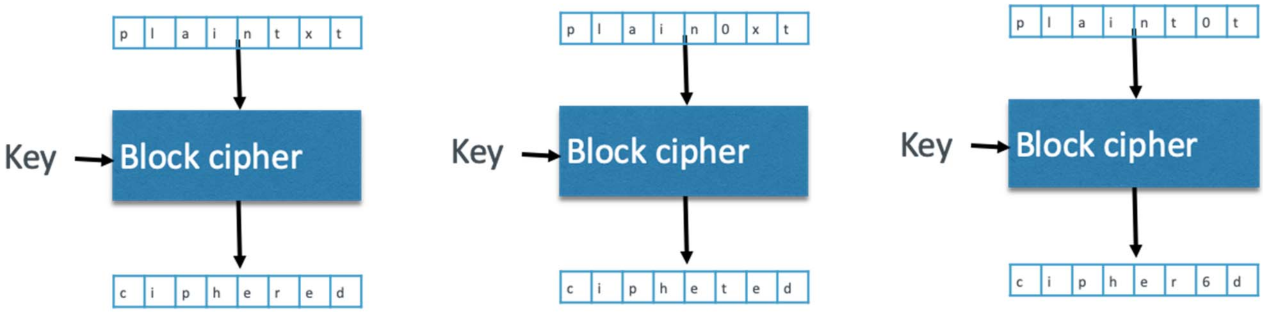

Figure A.2: ECB mode showing application of unmodified key at each instance of encryption

The consequence of ECB is that if two blocks of plain text are identical the resulting blocks of ciphertext are similarly identical. If the plain text is PII then when using ECB mode the resulting cipher text is directly mappable to the PII and every time the PII is encrypted the same ciphered representation of the PII will exist.

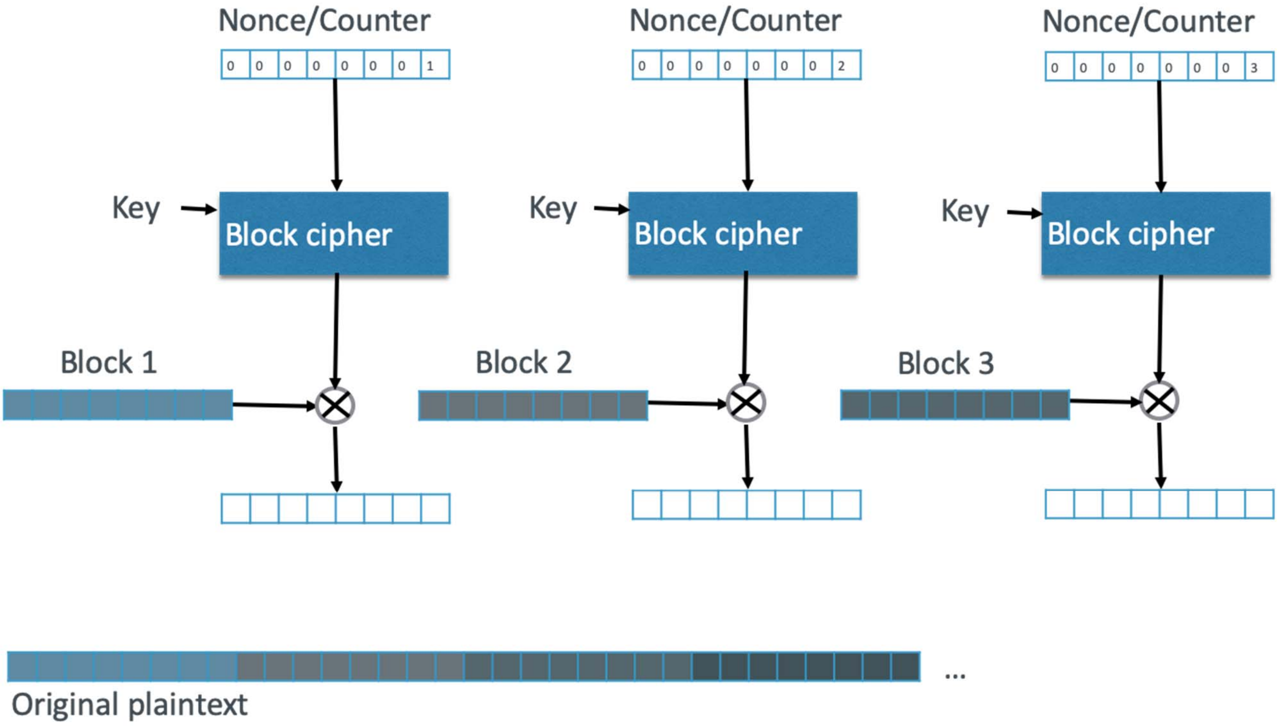

Figure A.3: CTR mode showing use of counter to modify key at each instance of encryption

In contrast to use of the ECB mode, use of the CBC or CTR mode results in different cipher text being produced for each instance of encrypting the same plain text block. This breaks the direct link between the PII in the plain text and the resulting cipher text.

A.6.2 Quantum computing safety

Systems based on hard problems such as the factorization of large numbers into prime number pairs are vulnerable to attack by Shor's algorithm on a quantum computer. Whilst the timetable for availability of a viable quantum computer is not known, due consideration by designers of cryptographic agility (see clause A.7) to ensure that data protected by an "at risk" algorithm can be migrated to a Quantum Safe Cryptographic facility. The text of ETSI EG 203 310 [i.1] and ETSI GR QSC 004 [i.2] apply to the risks of Quantum computing on cryptographically protected assets.

A.7 Cryptographic agility

All cryptographic schemes have a certain likelihood of being broken or weakened over time. In order to give assurance that protection mechanisms remain valid systems should be designed in such a way that algorithms and any required parameters can be updated over time. The general term for this is crypto-agility.

History

Document history		
V1.1.1	August 2020	Publication