ETSI GS INS 002 V1.1.1 (2010-09)

Group Specification

Identity and Access Management for Networks and Services
Distributed Access Control for Telecommunications
Use Cases and Requirements

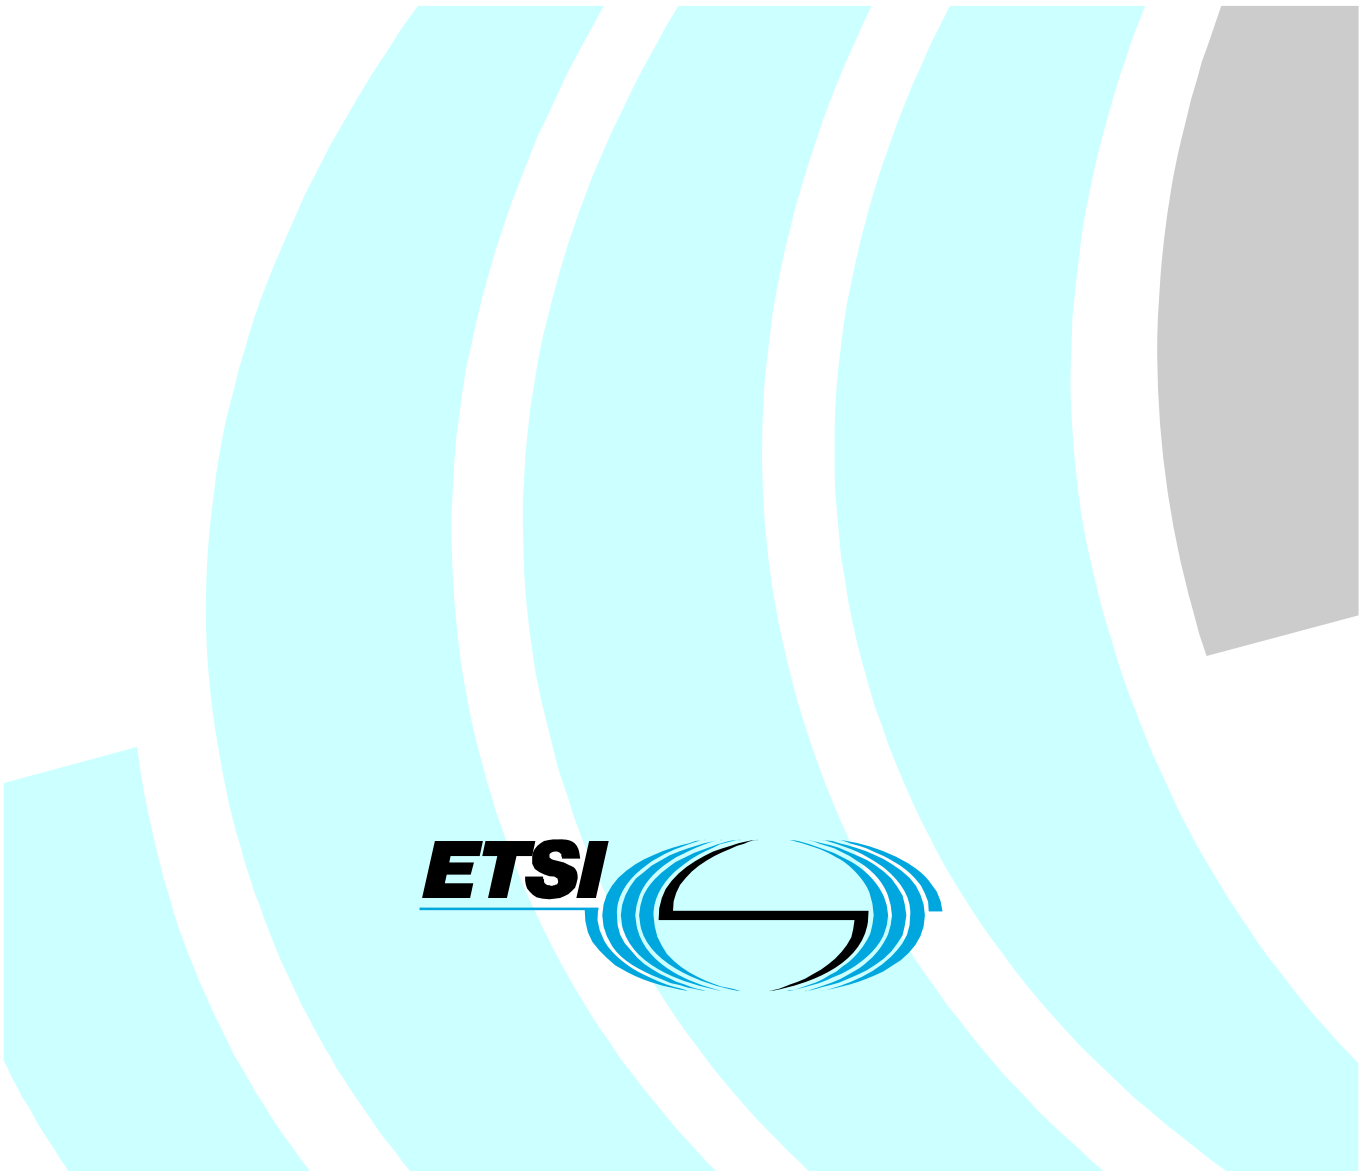

Reference
DGS/INS-002

Keywords
access, control, ID, management, network, service

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

Individual copies of the present document can be downloaded from:
http://www.etsi.org

The present document may be made available in more than one electronic version or in print. In any case of existing or perceived difference in contents between such versions, the reference version is the Portable Document Format (PDF). In case of dispute, the reference shall be the printing on ETSI printers of the PDF version kept on a specific network drive within ETSI Secretariat.

Users of the present document should be aware that the document may be subject to revision or change of status. Information on the current status of this and other ETSI documents is available at
http://portal.etsi.org/tb/status/status.asp

If you find errors in the present document, please send your comment to one of the following services:
http://portal.etsi.org/chaircor/ETSI_support.asp

Copyright Notification

Contents

Intellectual Property Rights

IPRs essential or potentially essential to the present document 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 (http://webapp.etsi.org/IPR/home.asp).

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.

Foreword

This Group Specification (GS) has been produced by ETSI Industry Specification (ISG) Identity and access management for Networks and Services (INS).

Introduction

Service and network providers need to restrict access to their functions in order to efficiently charge, protect critical systems and offer personalization. While historically this has been the case for many years, a new type of access control surrounding the user and its data becomes paramount in this day and age. Users are targeted by many different services, not all of them friendly, and require mechanisms to protect their data and information. In addition, the more social services are available, the more information about them is available and the harder it is to ensure that users' sensitive data would not be easily subject to theft and misuse.

In the present document we analyse not only the requirements for access control related to identity management but also bring this question one step further in considering that providers need to cooperate in order to enforce all the policies related to that user's data. This cooperation can be achieved either by exchanging data about the user or the context of the request, sharing policies or, in the case we will evaluate in this document, sharing the decision.

In the first part of the present document a summary of some of the activities around access control languages and mechanisms can be found. The second part of the document presents those use cases which we consider present new questions which are not yet addressed by other standardization activities. Finally, the third part of the document introduces a set of requirements extracted from the use cases.

1 Scope

The present document will provide requirements on the use and application of distributed policy management, decision and enforcement in a hybrid environment (operator and services domains).

2 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 reference document (including any amendments) applies.

Referenced documents which are not found to be publicly available in the expected location might be found at http://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.

2.1 Normative references

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

Not applicable.

2.2 Informative references

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] H. Schulzrinne, H. Tschofenig, J. Morris, J. Cuellar, J. Polk, and J. Rosenberg: "A Document Format for Expressing Privacy Preferences for Location Information", Nov 2003 Feb 2006, IETF draft (draft-ietf-geopriv-policy-08.txt).

NOTE: See http://tools.ietf.org/wg/geopriv/.

[i.2] H. Schulzrinne, H. Tschofenig, J. Morris, J. Cuellar, J. Polk, and J. Rosenberg: "Common Policy: A Document Format for Expressing Privacy Preferences", Feb 2004 Aug 2006, IETF draft (draft-ietf-geopriv-common-policy-11.txt).

NOTE: See http://tools.ietf.org/wg/geopriv/.

[i.3] H. Tschofenig, H. Schulzrinne, A. Newton, J. Peterson, A Mankin, The IETF Geopriv and presence architecture focusing on location privacy, Position paper at W3C Workshop on Languages for Privacy Policy Negotiation and Semantics-Driven Enforcement, Ispra, Italy, 2006.

[i.4] IETF RFC 4745: "Common Policy: A Document Format for Expressing Privacy Preferences".

NOTE: See http://www.rfc-editor.org/rfc/rfc4745.txt.

[i.5] T. Moses, eXtensible Access Control Markup Language (XACML) Version 2.0 OASIS Standard, Entrust Inc., 1 Feb 2005.

NOTE: See http://docs.oasis-open.org/xacml/2.0/access_control-xacml-2.0-core-spec-os.pdf.

[i.6] IETF RFC 2753: "A Framework for Policy-based Admission Control".

NOTE: See http://www.ietf.org/rfc/rfc2753.txt.

[i.7] ISO/IEC 10181-3 (1966): "Information technology - Open Systems Interconnection -- Security frameworks for open systems: Access control framework".

[i.8] Moses, T., ed., OASIS Privacy policy profile of XACML v2.0, OASIS Standard 1 Feb 2005.

NOTE: See http://docs.oasis-open.org/xacml/2.0/access_control-xacml-2.0-privacy_profile-spec-os.pdf.

[i.9] T.Moses et al.,"XACML Profile for Web Services", OASIS TC Working Draft, September 29th, 2003.

NOTE: See www.oasis-open.org/committees/download.php/3661/draft-xacml-wspl-04.pdf.

[i.10] Anderson, A., ed., Core and hierarchical role based access control (RBAC) profile of XACML v2.0; OASIS Standard, February 1, 2005.

NOTE: See http://docs.oasis-open.org/xacml/2.0/access_control-xacml-2.0-rbac-profile1-spec-os.pdf.

[i.11] Erik Rissanen et al., "XACML 3.0 administrative policy", working draft 07, 3 November 2008.

NOTE: See http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=xacml.

[i.12] IBM, Enterprise Privacy Authorization Language (EPAL), Version 1.2, 2003.

NOTE: See http://www.w3.org/Submission/2003/SUBM-EPAL-20031110/.

[i.13] A. H. Anderson. A comparison of two 5, New York, NY, USA, 2006. ACM Press.

[i.14] IETF RFC 3261: "SIP: Session Initiation Protocol".

3 Abbreviations

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

3GPP	Third Generation Partnership Project
AF	Application Function
BGF	Border Gateway Function
BGS	Border Gateway Services
BTF	Basic Transport Functions
CPN	Customer Premises Network
EPAL	Enterprise Privacy Authorization Language
FMC	Fixed and Mobile Converged
GBR	Guaranteed Bit Rate
GGSN	Gateway GPRS Service Node
GPRS	General Package Radio Service
IdM	Identity Management
IdP	Identity Provider
IP	Internet Protocol
ITU-T	International Telecommunication Union
MBR	Maximum Bit Rate
NACF	Network Attachment Control Functions
NAPT	Network Address and Port Translation
NASS	Network Attachment Sub-System
NAT	Network Address Translation
NO	Network Operator
OASIS	Organization for the Advancement of Structured Information Standards
OCS	On-line Charging System
PAP	Policy Administration Point
PCC	Policy Control and Charging
PCEF	Policy and Charging Enforcement Function
PCRF	Policy Control and Charging Rule Function
PD-FE	Policy Decision Function Entity
PDG	Packet Data Gateway
PDN	Packet Data Network
PDP	Policy Decision Point
PE-FE	Policy Enforcement Functional Entity

PEP	Policy Enforcement Point
PIP	Policy Information Point
QoS	Quality of Service
RACF	Resource and Admission Control Functions
RACS	Resource and Admission Control Subsystem
RBAC	Role Based Access Control
RCEF	Resource Control Enforcement Function
SaaS	Software as a Service
SBP	Service-Based Policy control
SCF	Service Control Function
SDO	Standards Development Organization
SDP	Service Delivery Platform
SIP	Session Initiation Protocol
SLA	Service Level Agreement
SPDF	Service Policy Decision Function
SPR	Subscription Profile Repository
TRC-FE	Transport Resource Control Functional Entity
UE	User Equipment
WiFi	Wireless Fidelity
WLAN	Wireless LAN
WS	Web Service
XACML	eXtensible Access Control Markup Language
XML	eXtended Mark up Language

4 Current Landscape

The need for Identity Management has surpassed the enterprise and web provider's world. Today, the need for Identity Management is present whenever the user needs to login or the provider needs information about the user. Information, authentication and authorization should be consistent and act as the glue between the different applications the user interacts with.

Access Control is one important aspect in the user's control of his/her identity. In a Distributed Identity Management Platform the user should be able to deal with various services, specify the preferences regarding the information revealed, coordinate between different hierarchical entities of an Identity Management Platform.

Access policies have the tendency to grow more complex over time. A policy may also depend on (generic) privacy policies which must be enforced by organizations due to legal regulations. Enterprises have complex policies instantiated by different divisions. In case of SaaS, two major reasons come into play: on the one hand these services could be composed from various elements, thus access to the service as a whole depends on access privileges from all those providers. On the other hand, the SaaS domain can be outside of the enterprise domain. Both these reasons imply that a monolithic policy definition is impractical. The key point of Distributed Access Control is to introduce an abstraction or modularization which splits the policies into building blocks. Through these building blocks, specific aspects (e.g. generic data privacy) could be specified. As each policy can have its own notion of subject, resources and action, an important aspect is the mapping of these definitions between different sets of policies.

4.1 General Access Control Frameworks

4.1.1 IETF Geopriv Working Group Policies Frameworks

The IETF Geographic Location/Privacy (Geopriv) working group has defined a protocol that carries Geopriv location objects (i.e. presence and location information). In term of privacy protection, the working group has published two frameworks for authorization policies controlling access to application-specific data, especially location and presence information; Common-Policy- [i.1] and Geopriv-Policy-Framework [i.2].

Whereas the Common-Policy, defines the basic rule structure using conditions, actions, transformations and provides a resolution mechanism that considers the fact that authorization policies might be used and evaluated in a distributed fashion, the Geopriv-Policy specifies several extensions of Common-Policy (i.e. with location-specific authorization policies with respect to conditions among others, or with presence authorization rules to perform authorization decisions for a presence based system) (see [i.3]).

In February 2007, the efforts of the Geopriv working group were standardized as RFC 4745 [i.4]. RFC 4745 [i.4] combines location- and presence-specific authorization aspects and uses XML to specify the language in which the privacy policy rules are to be represented.

The main purpose of the proposed policy framework is to protect the access on location and presence information. A policy consists of several rules. Each of these rules again consists of three parts: conditions, actions and transformations. A rule set defines conditions for data access and potential transformations/actions that should take place with the data. Conditions specify who is allowed to access the data when and under which conditions; e.g. whether the requested location information can be sent to the receiver or not. The transformation part of a rule describes optional changes (i.e. reduction of the location information's precision) of the requested location information before sending it to a requestor.

As identity, information serves any communication identifier, e.g. SIP address, mail addresses, etc. Validity restrictions can limit the applicability of a rule to a specific time period. Finally, it is possible to restrict the access by so called spheres, e.g. 'home', 'work', of the data owner. How this information is gathered is up to application specific realization. The modification of rules is controlled by so called metapolicies, which describe who is allowed to insert, update or delete a particular rule.

The RFC 4745 [i.4] common policy framework is designed to fulfil the needs of mobile application with requirement of being efficient and can be extended to other application domains.

4.1.2 eXtensible Access Control Markup Language

The eXtensible Access Control Markup Language (XACML) [i.5] is a widely adopted standard language for expressing both privacy and security policies in a machine-readable format (i.e. XML), developed and ratified by OASIS. Basically, XACML provides the syntax for a policy language and the semantics for processing those policies. It has, in accordance with RFC 2753 [i.6] and ISO 10181-3 [i.7], an abstract data-flow model and language model.

XACML policy language model, has three key components; rules that are the most elementary unit of a policy, they contain much complex logic aiming to make XACML expressive and fit to handle different policy requirements; policy that are combination of rules including a target, some obligations and an a rule-combining algorithm-identifier; and policy-set composed of a set of policies and several obligations (see [i.5]).

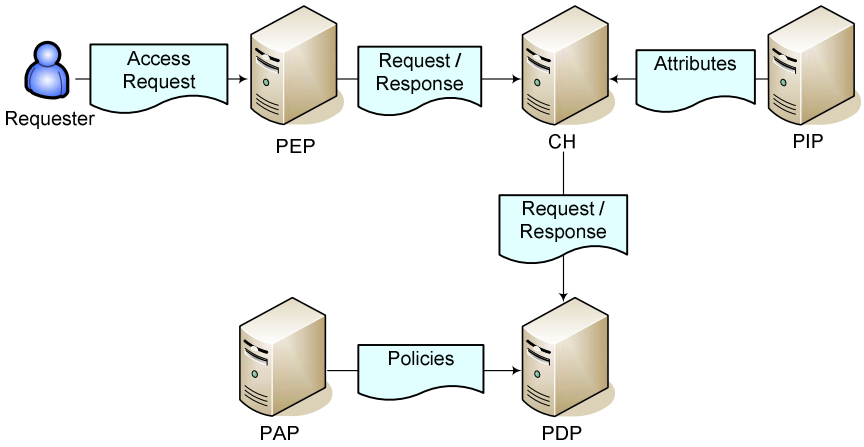

Figure 1: XACML Architecture

The XACML data-flow model shown in Figure 1 includes request and response formats to query policies, as well as semantics for determining validity of policies to the requests. Both request and response formats are seen as interface between two other abstract components: Policy Decision Point (PDP) and Policy Enforcement Point (PEP). All access requests for a protected data go through the PEP. The incoming access request together with other information (e.g. a description of which action is to be performed on the protected data) are condensed by the PEP and send to the PDP as request for an authorization decision. PDP then chooses the relevant policies to the decision request, along with any additional information (e.g. obligations; actions PEP must tackle along with enforcing the authorization decision) in some cases, evaluates the policies, and returns an authorization decision to the PEP. The resulted authorization decision is either "allow access", "deny access", or "no policies apply to this access request". Based on the resulted authorization decision, the PEP either grants or denies access to the requested data. Sometimes, neither "allow access" nor "deny access" is result from the policies evaluation. An error is returned instead if the policies could not be evaluated or retrieved.

The XACML 2.0 Standard defines several profiles, in order to tackle environments or technologies specifically demands and requirements e.g. management of attributes in web services or expressing policies that use role based access control. Another profile considers policy specification to archive privacy protection [i.8] or adopt parts of the RBAC standard to XACML [i.10].

The Privacy policy profile of XACML was defined in the version 2 of the XACML Standard in order to aid interoperability in the specification of privacy policies. It defines standard XACML Attribute identifiers for expressing the purpose for which data is collected and the purpose for which data is being accessed. It also defines a standard rule for requiring that the purpose for which data is collected must be consistent with the purpose for which data is being accessed.

Another extension is the Web Services Profile of XACML (WS-XACML) [i.9]) that defines one Web Services Policy Assertion targeting privacy concerns: The XACMLPrivacyAssertion. The XACMLPrivacyAssertion supports policy agreement, privacy/confidentiality requirements for PII, specific resources or parts of XML documents, data retention limits. Moreover, it fits into WSPolicy and can use P3P inside. XACMLPrivacyAssertion may be used to define privacy and confidentiality requirements like "You must not release my personal information to any 3rd party", or "You may keep my information less than 60 days."

A new XACML version, v3.0, is in preparation and will add generic attribute categories for administrative policy (i.e. who is allowed to write policies about what) and delegation, extend and integrate current XACML assertion types (e.g. the XACMLPrivacyAssertion) (see [i.11]).

4.1.3 Enterprise Privacy Authorization Language (EPAL)

The Enterprise Privacy Authorization Language (EPAL) is a formal language to express fine-grained enterprise privacy policies, ratified by the W3C consortium [i.12]. At its core, EPAL focus on the primary privacy authorization while abstracting data models and user-authentication from all deployment details.

In EPAL, privacy policy is defined as lists of hierarchical sorted data-, user-categories, and purposes and sets of privacy related actions, obligations, and conditions. The user-categories define entities (users/groups) that use collected data (e.g. travel expense department or tax auditor) while the data-categories set different categories of collected data that are handled in different ways from a privacy perspective (e.g. medical-record vs. contact-data). Purposes and actions define the service for which data is used (e.g. processing a travel expense reimbursement or auditing purposes), and how this data is used (e.g. disclose vs. read) respectively. Finally, obligations and conditions set restrictions in regard of actions (e.g. delete after 30 days or get consent) or context evaluation (e.g. "the user-category must be an adult" or "the user-category must be the primary care physician of the data-subject").

Enterprise privacy languages, EPAL particularly, are mostly used for company internal purposes and they are more fine-grained than the web privacy policy languages (i.e. XACML).

However, EPAL was not designed to be an all-round solution. For example, the manipulation of EPAL by data subjects in order to set their preferences (e.g. providing information that enables the creation of user-preferences forms based on the information given in an EPAL policy). In addition, Anderson [i.13] has shown that EPAL provides a subset of the functionality that can be provided by XACML.

4.2 Access Control in Telecommunications

Policy-based resource control provides the network with the intelligence required to manage transport network resources and adapt transparently to the different needs of running services and applications in aspects such as Quality of Service (QoS), authentication, authorization, accounting and charging. Policy-based resource control functions ensure that intended QoS is actually supported at the transport plane. This functionality is considered crucial for telecom operators in order to deliver 4-play services (voice, data, video and mobility) over Fixed and Mobile Converged (FMC) networks in a sustainable way. As such, different Standards Development Organizations (SDOs) have specified policy-based resource management functionalities.

In the field of standardization related to access control in telecommunications, the Telecommunication Standardization Sector of the International Telecommunication Union (ITU-T) has published the Resource and Admission Control Functions (RACF) specification, intended for general networks. The European Telecommunications Standards Institute (ETSI) has defined the Resource and Admission Control Subsystem (RACS), which targets fixed and mobile converged space, while the Third Generation Partnership Project (3GPP) has developed the Policy Control and Charging (PCC) Framework, specifically aimed for the mobile network environment that not only touches upon admission and resource management but also addresses charging. Although there are some differences between the various SDO's architectures (nomenclature, functional elements, research scope, and so on), the core concept remains the same, which is the clear separation between control and transport, between functionality and access technologies. In Figure 2 a small overview of all functional elements associated with access/resource control in the various mentioned SDO's network architectures is presented.

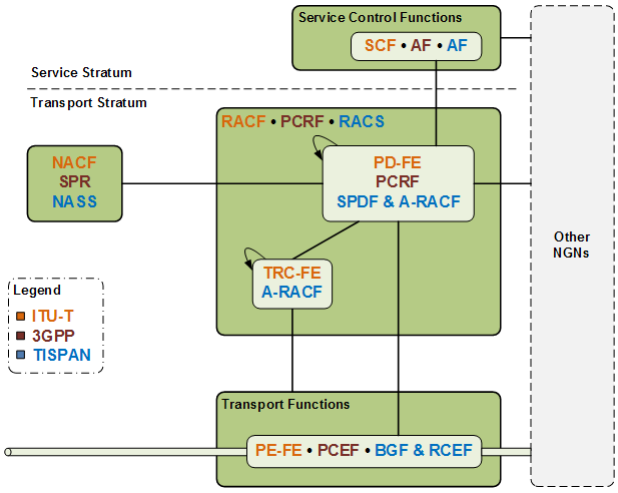

Figure 2: Policy-related architecture elements in ITU-T/3GPP/ETSI NGNs

With regards to QoS, all the three leading standardization organizations mentioned centralize their activity on control mechanisms for QoS guarantee, but do not address specific QoS technologies. They do have a unifying theme in their architectures, which is to implement resource and admission control function between service control and access/bearer layer in order to exercise control over resources of access/bearer nodes. The main aim is to shield specific technologies and topologies of the access/bearer layer from the service control layer. After receiving a service relevant QoS requirement from service control layer, resource control function combines it with the admission control strategies and network topologies, converts the service relevant QoS requirement into the IP QoS requirement, and then sends those requirements to related access nodes, bearer nodes and service gateway nodes. These nodes will then implement requested QoS in accordance with the messages received from the resource control elements.

4.2.1 3GPP Policy Control and Charging (PCC)

The policy control and charging (PCC) architecture allows operators to perform service based QoS policy and flow based charging control. PPC works on a service data flow level, and, as the name implies, provides the functions for policy and charging control as well as event reporting for service data flows. Its functionality encompasses two main functions:

- Flow Based Charging: including charging control and online credit control.
- Policy control including gating control and QoS control, among others.

Service flows are described by an aggregate set of packet flows characterized by identical source and destination IP address and port numbers. The PCC binds transport and service information in such a way that charging and policy are tied together to target heterogeneous transport networks. In the following paragraphs, a small description of each functional element and their role in 3GPP's PCC architecture is presented.

4.2.1.1 Application Function (AF)

The AF offers applications that require dynamic policy and/or charging control over the IP access network user plane behaviour. The AF communicates with the PCRF to transfer dynamic session information, required for PCRF decisions as well as to receive network specific information and notifications about bearer level events. One example of an AF is the P-CSCF in IMS. The PCRF may reject the service information request sent by the AF, and also inform it of service information that it would accept. In that case, the AF rejects the service establishment towards the UE and if possible the AF forwards the service information to the UE that the PCRF would accept. An AF may able to contact multiple PCRFs, choosing the appropriate one based on either the users' IP address and/or a UE (User Equipment) identity that it already is aware of.

4.2.1.2 Subscription Profile Repository (SPR)

The SPR logical entity contains all subscriber/subscription related information needed for subscription-based policies and network bearer level PCC rules by the PCRF. The SPR may be combined with or distributed across other databases in the operator's network, but those functional elements and their requirements for the SPR are not addressed.

4.2.1.3 Policy Control and Charging Rule Function (PCRF)

The PCRF is responsible for policy control decisions, flow based charging control, and monitoring capabilities (in regards to user plane traffic). It provides network control regarding the service data flow detection, gating, QoS and flow based charging (except credit management) towards the PCEF. Before accepting service information from the AF, the PCRF applies the security procedures, as required by the operator and crosschecks the service information provided by the AF to see if it is consistent with operator defined policy rules. It is also responsible for determining how a certain service data flow shall be treated in the PCEF, and ensure that the PCEF user plane traffic mapping and treatment is in accordance with the user's subscription profile.

The service information provided by the AF shall be used to derive the QoS for the service. As mentioned previously, the PCRF may reject the request received from the AF when the service information is not consistent with either the related subscription information or the operator defined policy rules and as a result the PCRF shall indicate that this in the response to the AF, along with the service information that can be accepted by the PCRF (e.g. the acceptable bandwidth).

To calculate the proper QoS authorization (QoS class identifier, bitrates), the PCRF uses the service information received from the AF (e.g. SDP information or other available application information, for both session based and non-session based services) and/or the subscription information received from the SPR.

4.2.1.4 Policy and Charging Enforcement Function (PCEF)

This functional entity is located at the Gateway (e.g. GGSN in the GPRS case, and PDG in the WLAN case). It provides service data flow detection, user plane traffic handling, triggering control plane session management (where possible), QoS handling, and service data flow measurement as well as online and offline charging interactions.

A PCEF ensures that an IP packet, which is discarded as a result from policy enforcement or flow based charging, is neither reported for offline charging nor cause credit consumption for online charging. The PCEF is enforces the policy control as indicated by the PCRF in two different ways:

- Gate enforcement: The PCEF only allows a service data flow to pass through it if and only if the corresponding gate.

- QoS enforcement: The PCEF is able to convert a QoS class identifier value to specific access network QoS attribute values and vice-versa. It is also in charge of PCC rule QoS enforcement, in which it enforces an authorized QoS of a service data flow according to active PCC rules. Finally, the PCEF controls the QoS that is provided to a combined set of service data flows. The policy enforcement function ensures that the resources which can be used by an authorized set of service data flows are within the "authorized resources". The authorized QoS provides an upper bound on the resources that can be reserved (GBR) or allocated (MBR) for the access network bearer.

Enforcement of charging control policies is another function within the PCEF realm. For a service data flow (defined by an active PCC rule) that is subject to charging control, the PCEF will allow the service data flow to pass through the PCEF if and only if there is a corresponding active PCC rule with and, for online charging, the OCS has authorized credit for the charging key. The PCEF may let a service data flow pass through the PCEF during the course of the credit re-authorization procedure.

In short, for any service data flow (defined by an active PCC rule) that is subject to both Policy Control and Charging Control, the PCEF will only allow it to pass through if and only if the right conditions from both policy control and charging control happen, meaning that the corresponding gate is open and, in case of online charging, the OCS has authorized credit for its charging key.

A PCEF may be served by one or more PCRF nodes. The PCEF shall contact the appropriate PCRF based on the packet data network (PDN) connected to, together with, a UE identity information (if available, and which may be dependent on the access network used).

4.2.2 ETSI TISPAN Resource and Admission Control Sub-systems (RACS)

ETSI TISPAN architectural element related to access control can be surmised by 3 sub-systems: the Network Attachment Sub-System (NASS), the Resource Admission Control Sub-System (RACS) and the Transporting Processing Functions.

The **Network Attachment Sub-System (NASS)** provides registration at access level and initialization of User Equipment (UE) for accessing to the TISPAN NGN services. The NASS provides network level identification and authentication (based on user profile), manages the IP address space of the Access Network and authenticates access sessions. The NASS also announces the contact point of the TISPAN NGN Service/Applications Subsystems to the UE. Network attachment through NASS is based on implicit or explicit user identity and authentication credentials stored in the NASS.

The **Resource Admission Control Sub-System (RACS)** is responsible for elements such as policy control, resource reservation and admission control. In addition, it also supports core Border Gateway Services (BGS) including Network Address Translator (NAT) mechanisms. The RACS provides policy based transport control services to applications, allowing for the request and reservation of transport resources from access and core transport networks within its coverage. By hiding the interaction between applications and transport resources, RACS ensures that applications do not need to be aware of the underlying transport networks. As the network system responsible for policy based transport control, RACS is able to perform admission control in order to evaluate those requests in the context of predefined policy rules provisioned by the network operator. It may then perform resource reservation provided the request passes the policy tests and appropriate resources are available in the transport network. In this way, RACS offers the operators the means to perform admission control, which may be followed by the installation of bearer service policy rules. In terms of session awareness, RACS is resource-reservation session aware but application session agnostic (i.e. it can support transport resource reservations for both session based and non-session based applications).

The **Transporting Processing Functions** include elements such as the Resource Control Enforcement Functions (RCEF) and the Border Gateway Function (BGF), whole role is to enforce the policies passed down by the RACS and also to provide network status information to the control plane, information which then may be used as input for certain policies.

In the following paragraphs a small description of each functional element and their role in TISPAN's RACS architecture is presented.

4.2.2.1 Application Function (AF)

The AF provides information to the SPDF to identify media flows to express the service expected from RACS and the bandwidth that needs to be authorized and allocated for those flows. Bandwidth requirements are complemented with class based service information indicating service expectations such as QoS characteristics. This class-based information may also capture predefined traffic characteristics. Resource priority requirements may also be supplied. The AF also indicates whether the media should be enabled (i.e. gate opened) when resources are allocated, or if that the gate should be opened later, after resources are committed.

The AF may be capable of operating in a mode of operation by means of which the AF request resources for media flows belonging to a single application session per resource request. The AF is capable of operating in any or all of the following modes of operation:

- the mode where a single resource reservation request per application session is issued by the AF;
- the mode of operation where multiple independent resource reservation requests per application session are issued either from a single or multiple AFs, where each independent request is intended to reserve a different set of resources within the network.

The AF is entitled to use Subscriber-Id and/or an IP address to identify to RACS the resource being requested. The decision of what information is provided to RACS depends on the type of application.

4.2.2.2 Service Policy Decision Function (SPDF)

The SPDF is a Functional Entity that acts as a final Policy Decision Point for Service-Based Policy control (SBP) for each administrative domain it resides in and provides the application layer with a single point of contact, hides the topology of bottom network and specific access type, and supplies service-based policy control. In addition, SPDF selects local policy according to request from the Application Function (AF), maps the request into IP QoS parameters, and sends them to A-RACF and the Border Gateway Function (BGF) to request for corresponding resources (all according to service policy rules defied by the network operator).

4.2.2.3 Generic Resource and Admission Control Function (x-RACF)

The generic Resource Admission and Control Function (x-RACF) is a Functional Entity that acts as a Policy Decision Point (PDP) in terms of subscriber access admission control, as well as in terms of resource handling control. It receives request from SPDF. Based on the stored policies, admission control is realized by accepting or refusing the request for transport resources. The A-RACF obtains the network attachment and subscriber QoS list information from the Network Attachment Subsystem (NASS). Accordingly, available network resources can be assured with the network location information (as the physical node address of the access subscriber). Meanwhile, the subscriber QoS list information is referred to in the process of resource request.

Two functional specializations of the generic Resource Admission Control Function are defined: Access-RACF (A-RACF) and Core-RACF (C-RACF), which can be deployed in different network domains based on the operator's requirements. The main distinction between A-RACF and C-RACF is that the A-RACF checks the subscriber QoS profile that may be obtained from the NASS. The A-RACF is deployed in the access network domain, which may require the provisioning of the transport resources on a per subscriber basis; On the other hand the C-RACF does not check the subscriber QoS profile. The C-RACF is deployed in the core transport network domain, which may not provision the transport resources on a per subscriber basis.

4.2.2.4 Border Gateway Function (BGF)

The BGF is a packet-to-packet gateway for user plane media traffic. The BGF performs both policy enforcement functions and NAT functions under the control of the SPDF in each of the network segments: access, aggregation and core.

4.2.2.5 Resource Control Enforcement Function (RCEF)

The Resource Control Enforcement Function (RCEF) performs policy enforcement functions for unicast and/or multicast after installation of traffic policies under the control of the x-RACF. Depending on the policy enforcement request, RCEF either enforces the policy autonomously (i.e. without involving other functional entities) or in conjunction with the BTF (e.g. trigger transport control actions).

The RCEF is managed by the RACS and is usually located in a Transport Network node Element, although it may exist in Access Network Domains as well as in Core Network Domains.

4.2.3 ITU-T Resource and Admission Control Functions (RACF)

The RACF provides an abstract view of transport network infrastructure to the Service Control Functions (SCF) and decouples the provision of services from the details of transport facilities such as network topology, connectivity, resource utilization and QoS mechanisms/technology, etc. The RACF interacts with the SCF and transport functions for a variety of applications (e.g. SIP-based call (b-IETF RFC 3261 [i.14]), video streaming, etc.) that require the control of NGN transport resource, including QoS control, NAPT/firewall control and NAT traversal.

The RACF executes policy-based transport resource control upon the request of the SCF, determines transport resource availability, makes admission decisions, and applies controls to transport functions for enforcing the policy decisions. The RACF interacts with transport functions for the purpose of controlling one or more of the following functions in the transport stratum: bandwidth reservation and allocation, packet filtering; traffic classification, marking, policing, and priority handling; network address and port translation; firewall.

The RACF takes into account the capabilities of transport networks and associated transport subscription information for subscribers in support of the transport resource control. The RACF interacts with network attachment control functions (NACF), including network access registration, authentication and authorization, parameter configuration, etc., for checking transport subscription information.

The RACF consists of two types of resource and admission control functional entities: the PD-FE (policy decision functional entity) and the TRC-FE (transport resource control functional entity). This decomposition of PD-FE and TRC-FE enables the RACF to support a variety of access and core networks (e.g. fixed and mobile access networks) within a general resource control framework. In the following paragraphs a small description of each functional element and their role in ITU-T RACF architecture is presented.

4.2.3.1 Service Control Function (SCF)

The SCF represents an abstract notion of the functional entities in the service stratum of NGN that request the QoS resource and admission control for media flows of a given service.

4.2.3.2 Policy Decision Function Entity (PD-FE)

The PD-FE provides a single contact point to the SCF and hides the details of transport network from the SCF. The PD-FE makes the final decision regarding network resource and admission control based on network policy rules, SLAs, service information provided by the SCF, transport subscription information provided by the NACF in access networks, and resource-based admission decision results provided by TRC-FE. The PD-FE controls the gates in the PE-FEs at a per flow level. The PD-FE consists of transport technology-independent resource control functions and is independent of the SCF as well. The policy rules used by PD-FE are service-based and are assumed to be provided by the network operators.

4.2.3.3 Network Attachment Control Functions (NACF)

The NASS equivalent of TISPAN ITU-T's NACF includes a collection of functional entities that provide a variety of functions for user access network management and configuration based on the user profiles. Among these features are: Dynamic provision of IP address and other user equipment configuration parameters; authentication of user access network, prior or during the IP address allocation procedure; authorization of user access network, based on user profiles (e.g. access transport subscription); access network configuration, based on user profiles; and location management.

4.2.3.4 Transport Resource Control Functional Entity (TRC-FE)

The TRC-FE deals with the diversity of underlying transport technologies and provides the resource-based admission control decision results to the PD-FE. The TRC-FE is service-independent and consists of transport technology-dependent resource control functions. The PD-FE requests the TRC-FE instances in the involved transport networks to detect and determine the requested QoS resource along the media flow path. The TRC-FE may collect and maintain the transport network topology and the transport resource status information and authorize resource admission control of a transport network based on network information such as topology and/or connectivity, network and element resource availability, as well as the transport subscription information in access networks.

4.2.3.5 Policy Enforcement Functional Entity (PE-FE)

The PE-FE (policy enforcement functional entity) in the transport stratum is a packet-to-packet gateway at the boundary of different packet networks and/or between the CPN and access network. It is the key injection node to enforce dynamic QoS and resource control, NAPT control and NAT traversal.

5 Use Cases

5.1 UC1: Software as a Service

5.1.1 Description

In a first use case, we assume that an employee of a specific company is utilizing a Software as a Service (SaaS) platform. While the core functionality is provided by the provider of the SaaS platform, additional services are provided by others. The employee's company wants to limit the access to this SaaS service not only to a particular set of its work force but also keep control of the data used for this service in order to avoid unauthorized access and misused of such sensitive data.

The related policies of the company already contain some confidential information, thus, they should not be distributed instead a distributed evaluation of the policies is seeked. In addition, some obligations have to be fulfilled by the SaaS platform due to the access request. The SaaS can integrate the different policies through generic distributed policies, ensuring that together with its own policies, the policies of the company are respected and privileges to use the external services are actually available.

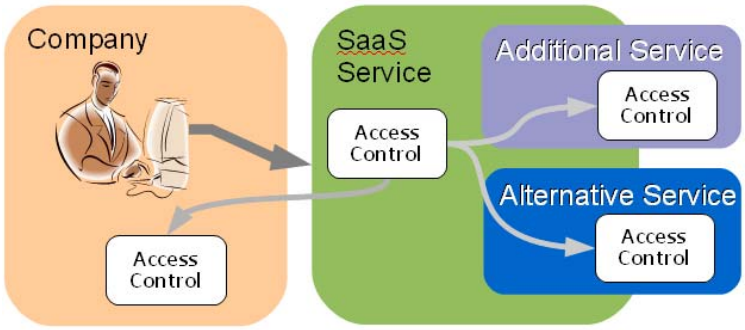

Figure 3: Use Case SaaS

This use case is sketched in Figure 3. As an important aspect neither the internal policies of the company nor those of the additional services have to be known at the SaaS service. The obligations send to the SaaS service based on the underlying policies have to be understood and adhered at the actual SaaS service.

In this use case the role of SaaS Service Provider may be performed by multiple types of entities. While it can be a dedicated SaaS Provider, an operator is also in a good position to provide such services.

5.1.2 Actors

- User, who wants to utilize a SaaS service for his job activities.
- Company, has outsourced a service, but wants to keep control on the processes.
- SaaS service provider, handling the service by utilizing other service providers.
- Additional Service and Alternative Service providers, offering their services to be incorporated into a larger SaaS environment.

5.1.2.1 Actors specific Issues

- User:
  - Wants to use the SaaS service.
- Company:
  - Wants its employees to use the SaaS service under well defined conditions.
  - Does not want to share the actual policies for accessing the service with any 3rd party entity.
  - Wants its obligations related to the policies being enforced.
- SaaS Provider:
  - Provides a service to the customer.
  - Coordinates the access to the service with company and other services.
  - Has to know and enforce all obligations.
- Additional service and Alternative service providers:
  - Provide services for larger environments.
  - Have their own policies regarding utilization of their services.
  - Do not want to share the actual policies for accessing the service with any 3rd party entity.
  - Want their obligations related to the actual policies being enforced.

5.1.2.2 Actors specific benefits

- User:
  - Utilizes the SaaS service as if it is inside the company.
- Company:
  - Outsources the service but still can enforce its policies.
  - Keeps the actual policies confidential.
  - Ensured that the obligations related to the policies are enforced.
- SaaS Provider:
  - Enjoys business advantages for being able to provide the service to multiple companies.
  - Can provide a feature for managing access according to the company/customer policies.
- Additional service and Alternative service providers:
  - Enjoys business advantages for being able to provide services for larger environments.
  - Keep the actual policies confidential.
  - Ensured that the obligations related to the policies are enforced.

5.1.3 Pre-Conditions

- SaaS Provider has exchanged the obligation enforceable at his side with the company and the additional/alternative services.
- User is authenticated.

- Endpoints to send authorization request to the Company and additional/alternative services which are known and available to the SaaS service.

5.1.4 Post-Condition

- User has utilized the SaaS service according to the policies of the SaaS provider, company and additional services.

5.1.5 Normal Flow

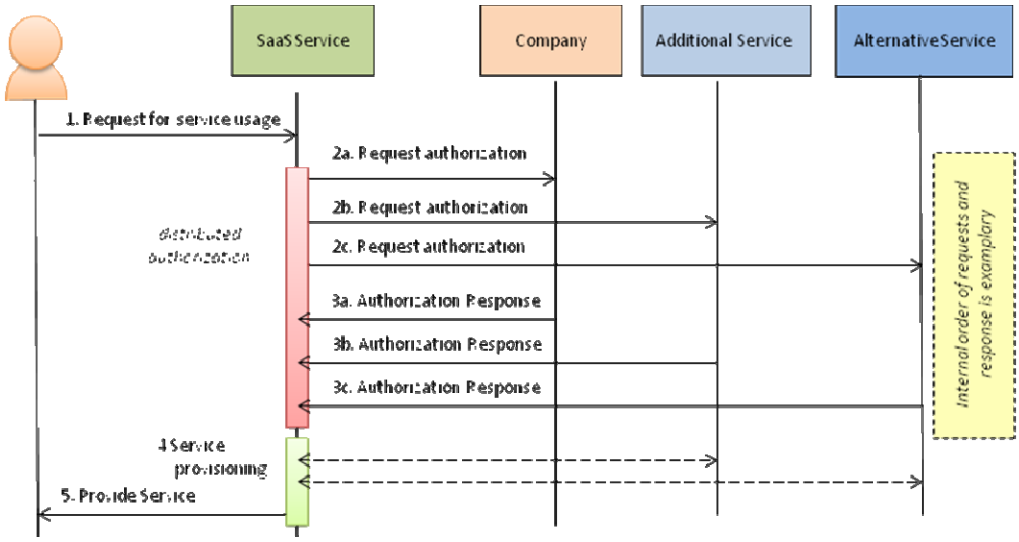

Figure 4: Flow of Use Case SaaS

1) Authenticated user sends a request to the service.

2) The SaaS service performs a distributed authorization evaluation based on the distributed policy sends authorization requests to:

   a) The company of the requesting user.

   b) Additional service.

   c) Alternative service.

   The actual order of this requests is exemplary and depends on the actual policy.

3) The related authorization engines send back their results. The order is arbitrary and depends on the actual deployment. Depending on this the final decision is reached.

4) The actual service is provisioned, the details on the depending messages send between SaaS and the other services are for simplicity reasons not shown in Figure 4.

5) Finally, the service is provided to the user.

In case the original request for a service is denied, the flow is quite similar. Step 4 is not executed and in Step 5 a denial message will be send back.

5.2 UC2: Enterprise Environment

5.2.1 Description

In this use case, it is assumed that a company has equipped its employees with mobile phones and uses a Presence Service of the related Telco operator to provide information on the employee for further coordination. The Presence Service might not only indicate general availability of a person but also the location or the proximity of other employees to indicate meetings etc. As a general Presence Service of the operator is available not only to the employees of the company the information provided by this service and to whom has to be guarded to avoid leakage of information. The Presence Service has to check to whom presence information are sent and whether any obligations such as obfuscation the information might be necessary. The actual decision could be taken only by the company based on internal information which should not be shared with the service provider. As an example, let's assume that the employee is visiting a customer in a particular town, while this customer might get detailed presence information to see that the employee has arrived on schedule at the airport etc. Other customers who are visited the next week do not get any detailed information at all. Instead it is just indicated that the employee is travelling or in a meeting.

Again, the policies have to be evaluated in a distributed fashion and the obligations send from the company to the Presence Service have to be well defined and enforced.

5.2.2 Actors

- Pete, employee of a company, publishing the presence information.
- Sara, subscribing to Pete's presence information.
- Presence Service providing the information to subscribers.
- Company, controlling the presence information of its employees.

5.2.2.1 Actors specific Issues

- Pete:
  - Publishes the presence and other information.
- Sara:
  - Subscribe to information on Pete.
- Presence Service:
  - Provides the presence information according to the policies of the related company.
- Company:
  - Controls the disclosure of Pete's presence information.

5.2.2.2 Actors specific Benefits

- Pete:
  - Utilizes the Presence Service.
- Sara:
  - Receives information via the Presence Service.

- Company:
  - Utilizes the Presence Service to assist its employees in carrying out their job responsibilities.
  - Keeps its actual policies confidential.
  - Ensures that obligations related to these policies are enforced.
- Presence Service:
  - Enjoys business advantages for being able to provide the service to multiple companies.
  - Can offer a feature to manage access according to the company/customer policies.

5.2.3 Pre-Conditions

- Presence Service has exchanged the obligation enforceable at his side with the Company.
- User are authenticated.
- Endpoint to send authorization request to is the Company which need to be known and available to the SaaS service.

5.2.4 Post-Conditions

- Users have utilized the Presence Service according to the policies of the SaaS provider, company and additional services.

5.2.5 Normal Flow

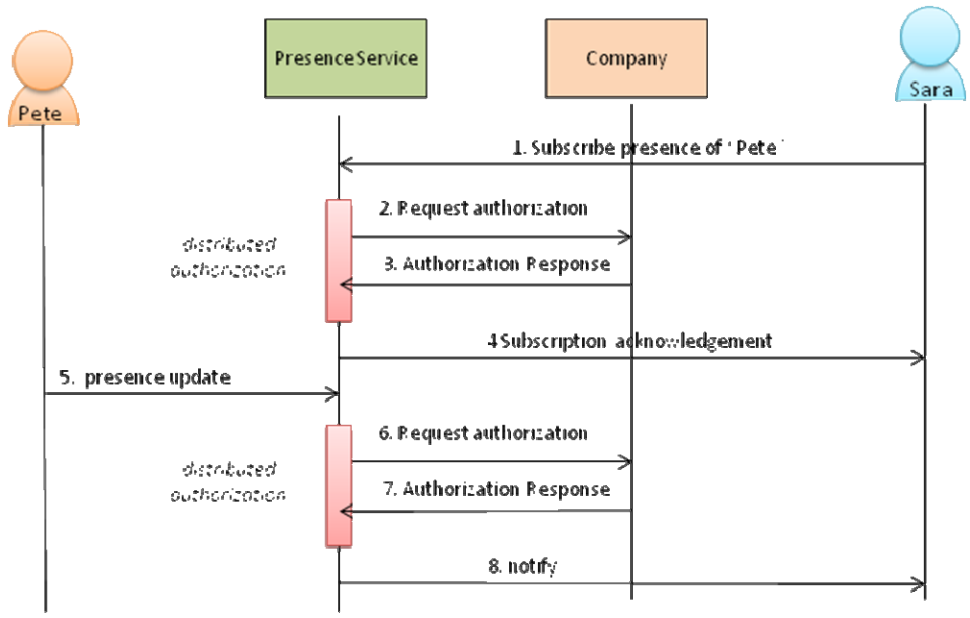

Figure 5: Normal Flow Use Case Enterprise Environment

1) Sara sends a subscription request for any presence information on user Pete.
2) The Presence Service evaluates the distributed policies and thus sends an authorization request to Pete's company.
3) The responsible entity evaluates the request and sends back a response. Depending on the results from the company the final decision regarding Sara's subscription is taken.
4) Assuming that the subscription was granted, the acknowledgment is send back to Sara.

5) Pete updates his presence information.

6) In order to find out to whom this particular presence update has to be send, the Presence Service evaluates the distributed policies and thus sends an authorization request to the company of the related user.

7) The responsible entity on Pete's company side evaluates the request and sends back a response. Depending on the results from the company the final decision is taken.

8) Assuming that the subscription was granted, the notification on the updated presence information is send to Sara.

In case the step 7 is a denial for the user Sara, obviously no message is send in step 8.

5.3 UC3: Roaming Network Access

5.3.1 Description

This use case demonstrates how a decision, on whether a service (in this case network access) can be granted or not, can be made based on multiple policy evaluations by different entities.

An Employee X is going on a business trip, and while at the Airport he wants to connect to the internet to perform company-related business. However he does not own an account with the Network Operator (NO) that provides WiFi at the Airport.

His company has, however, an agreement with the NO to provide internet access, although limited to certain sites, and only available to some of their employees. Employee X thus signals the NO that he would like to sign in with his company's credentials. He's redirected to his company's authentication portal along with relevant information from the NO concerning the access desired. Once Employee X credentials are validated by his company, several checks are done by this very same entity to determine whether this network access request by that specific employee is in line with company policies and whether or not it can be granted. After this evaluation is done, the result is sent back to the NO along with the authentication assertion and some Employee X attributes (e.g. his name). As the Airport WiFi can only be used by valid ticket carriers, the NO still has to verify if the Employee X possess one of these before the internet access can be granted. The NO, therefore, contacts the Airport information system in order to confirm that very same fact. Once all conditions are satisfied, network access is granted and later billed to Employee X's company.

5.3.2 Actors

- Employee X.
- Employee X's Company.
- Network Operator (WiFi Provider).
- Airport Information System.

5.3.2.1 Actors Specific Issues

- Employee X:
  - Desires internet access but does not have an account with the WiFi provider.
- Employee X's Company:
  - Stores data related to their employees.
  - Acts as an IdP for its employees.
  - Has defined specific policies regarding company-provided services employees have access to.
- Network Operator (WiFi Provider):
  - Controls the network access infrastructure.

- Has a trust relation with Employee X's company.
- Has a trust relation with airport information system.
- Acts as a service provider and enforcement point (regarding the decision on whether to grant network access or not).

- Airport Information System:
  - Stores data regarding travellers and corresponding flights.

5.3.2.2 Actor Specific Benefits

- Employee X:
  - Can access the internet without having to create yet another account with another provider.
- Employee X's Company:
  - Maintains control over employee authentication and authorization procedures (even when the services themselves are provided by 3rd parties).
  - Can provide internet and/or intranet access to any employee en masse (and in several locations) with a single agreement with one Network Operator.
- Network Operator:
  - Can leverage its installed infrastructure to provide added value services to specific audiences under controlled conditions.
  - Ensures network access control to its infrastructure.
  - Does not have the onus of managing users credentials or attributes.

5.3.3 Pre-conditions

- Network Operator and Employee X's Company are in the same circle of trust.
- Network Operator and Airport information system are in the same circle of trust.
- Employee X's company has a pre-defined set of policies regarding the service of network access for its employees.
- Employee X is provisioned at his company.

5.3.4 Post-conditions

- Employee X gains internet access only by using his company-given credentials and according to subset of policy common to all involved players.

5.3.5 Example Flow

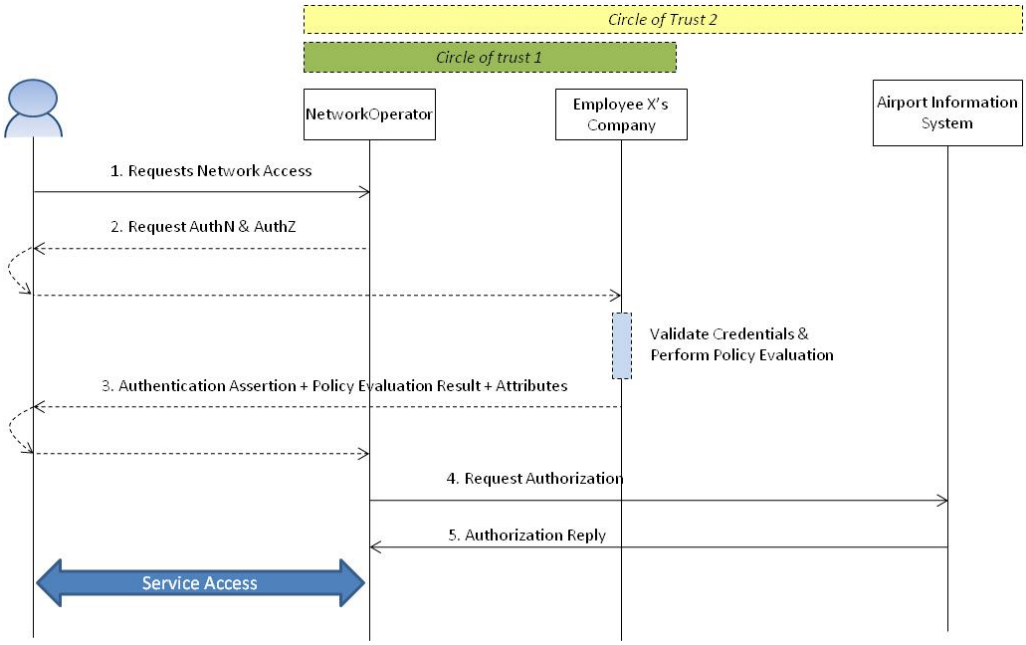

Figure 6: Roaming network access example flow

1) Employee X access the WiFi portal run by the Network Operator and signals the intention to authenticate using his corporate identity.

2) Employee is redirected to his company portal to perform authentication.

3) After a successful authentication the company verifies if this specific employee can make use of this service (network access). The result of these operations (authentication and authorization) is transmitted back to the Network Operator along with some requested (and agreed upon by both parties) attributes.

4) Network Operator asks the Airport information system if the employee is has a valid ticket.

5) Upon confirmation the Network Operator grants the employee access.

5.4 Summary Table of Use Cases

Table 1

Number	Summary
UC1	SaaS can integrate different policies through generic distributed policies, ensuring that together with its own policies, the policies of the company are enforced and privileges to use the external services are actually available.
UC2	In this use case it is assumed that a company has equipped its employees with mobile phones and utilizes a presence service of the related Telco operator to provide information on the employee for further coordination.
UC3	This use case demonstrates how a decision, on whether a service (in this case network access) can be granted or not, can be made based on multiple policy evaluations by different entities.

6 Requirements

This clause aims to explore requirements and exploit an overall IdM framework which provides consistent crucial services to both the network and the services that compose it. The architecture's core is an IdM framework which encompasses many different protocols while being conceptually sound and bringing convergence in core functions of these protocols. The IdM framework will then provide facilities to network and service functions to access identity related information such as authentication, attributes, access control, delegation, billing, etc. In clauses 6.1 to 6.5, the requirements to be met by the framework that were extracted from the aforementioned Use cases are listed.

6.1 General Access Control Framework Requirements

R1. In order to support transparency, there should be a mechanism for an entity to provide evidence that it needs certain information from the user and an interface for external auditing in terms of privacy policies and data processing.

We support transparency in relation to the Identity Management framework with two mechanisms: evidence that certain information is required by the service and support for external auditing mechanisms on the handling of identity data.

R2. Authentication, Integrity and non-repudiation should be enabled for all transactions.

Transactions among Users, Service Providers, Identity Brokers, Authentication Providers and Attribute Providers must be able to provide authenticity, integrity and non-repudiation **t**o avoid third parties to be able to access or modify identity data and to avoid entities to reject responsibility on their actions.

R3. Support of granular authorization.

The framework should support granular authorization which supports one or multiple transactions. Payment sometimes depends on how often an authorization event takes place e.g. for service bundling. Flexible payment methods are needed. Also, more real time access control is required.

6.1.1 Policy Management

R4. Authenticity, integrity and non-repudiation should exist between the different interacting entities.

R5. Users should have a simple mechanism to both set and assess the consequence of policies; even when these policies are set by an agent on behalf of the user.

The construction of policies over user data is a strenuous and complex activity for the typical user. Whether it is the user setting up his policies or an agent acting on his behalf, it is important that the user understands the practical results of these policies being applied to his data in the context of the affected services.

R6. Enable authorized personnel to audit the status and usage of the security mechanisms.

The auditing personnel should be able to prove that a certain transaction was made and provide information about the parties involved. In order to provide these assurances, security mechanisms should be combined with logging facilities while preserving the user's right to privacy.

R7. Availability of user's privacy preferences.

Although the authorization using user's attributes is spread widely in common federations, the usage of preferences is not mentioned. So, a user shall be given the possibility to define preferences that are stored like attributes and also transmitted within the federation. Those preferences can be taken by SPs to provide more comfort to the user.

R8. The framework must support dynamic management of subjects, policies and target at any time.

Policies added at runtime must be activated as soon as possible to enable the timely change policies. Access to policy management must be ensured at any time.

R9. The access control framework must support the delegation of rights.

Through delegation of rights, a user can be allow to act on behalf of another user, who may be the subscriber of a service.

R10. The possession of attributes must be unforgeable.

Valid proof of ownership of an attribute should only be generated by legitimate users.

6.1.2 Decision

To make a decision, the context around the request should be stated and both the context and subjects authenticated. Unless this requirement is met, the authorization decision will be made on non-secure data and may reflect the wrong decision.

R11. Authentication assertion and authentication context should be available for the authorization.

R12. If an authentication assertion could not be directly understood by the original requestor, a method to transform the assertion and related data should be provided by trusted entities.

The decision engine should support different types of resources and subjects. In addition, delegation cases should be handled seamlessly. When a decision is processed in this way, the associated policies become easier to define.

R13. The requestor is authenticated and is either a user, an application acting on behalf of a user, or a machine running an application and/or under the control of a particular user.

R14. The authorization for a particular type of access should be based on a request which includes related attribute information and the resource with related information.

R15. Authorization requests should be responded within a well defined time frame, or a default reaction should be enforced.

R16. Authorization responses may include addition obligations which have to be enforced as a reaction of the request independently whether the response was a denial or a permit.

When an access control decision is made, the policy may add additional commands to the enforcement point. These commands take the form of obligations which are part of the enforced decision.

6.1.3 Enforcement

In order to guarantee the correct enforcement point is being triggered, an authentication procedure needs to be in place. The authentication may occur using many different mechanisms which may be different in nature. The correct authentication level must be defined and provided.

R17. User agent should be able to authenticate to a mutually agreed authentication server.

R18. Different types of authentication technologies or protocols can be supported.

The authentication process itself may also be distributed. Especially when defining a cross domain decision, it is critical the authentication process can offload to another domain.

R19. Authentication request might be forwarded to another authentication server.

R20. An authentication server function should exist and should be able to create assertions about the user's identity.

R21. Consistent policy enforcement must be available on each layer of the framework.

Policy enforcement must be assured at each layer to prevent the circumvention of the access control mechanisms. Using access control mechanisms on different layers in the architecture implicates the policies to be consistent and may include that the policy must be translated to each layer (services, network, etc).

In UC1 SaaS, provider might be trusted to properly manage outsourced services and related data, but not authorized to use them. Therefore additional security solutions (e.g. encryption scheme) are required which support the reliable enforcement of access control policies on outsourced service/data.

NOTE: Outsourced resources might be encrypted according to an encryption obligation regulated by authorizations.

6.2 Distributed Access Control Requirements

R22. Establishment of Trust Relationship.

Since critical information is being transmitted within the federation, a trust relationship has to be established. Without an underlying trust relationship, on the one hand no one would allow providing private data and on the other hand no one would rely one received data to allow access to restricted domains, resources or services.

R23. No spread of security breaches.

Depending on the established trust relationship, information from federation members is highly trusted. Therefore, a successfully done attack on one federation participant may lead to the distribution of compromised/false information. There shall be no chance of infecting or attacking other entities within the federation with such information.

R24. Retrieval of attributes from several different Attribute Providers must be possible.

When leaving the academic sector, for which many federations were be designed, the way of storing attributes changed. Due to privacy issues, user attributes are not stored centralized but are spread over several attribute providers. There should be a way to collect those attributes from the several attribute providers without breaking user's privacy.

R25. The framework must support the combination of distributed or cascaded policies from different administrative entities.

To enhance privacy of an object, the framework must be able to derive decisions on policies at different locations.

6.2.1 Policy Management

R26. A central point collecting all the policies of different entities should be avoided.

R27. The Identity management framework must provide services and users the way to discover Identity Brokers (for Single Sign On/Single Log Out) and Attribute Providers (for attribute exchange), and obtain user's attributes from them under user's control, with using user's pseudonym or anonym.

Discovery functions will allow IdPs, data and functions to be dynamically discovered from a previously unknown domain while pseudonyms prevent unauthorized entities from linking the user to transaction. Discovery operations are performed under a strong privacy assumption.

R28. The policy-based access control framework should provide means for managing the overall policy life cycle, i.e. by providing functions for specifying, monitoring, enforcing and de/activating policies or providing mechanisms to guarantee the secrecy of policies (since sensitive information related to the policy can be deduced from the exchange between interacting entities even when the policy itself is not disclose).

6.2.2 Decision

R29. In a distributed environment authorization decisions may depend on decisions of other entities. The requesting entity is responsible for combining the results.

This step may be done recursively but mechanisms should be established to ensure it is finite.

R30. In case the final decision depends on multiple decisions by different entities all the obligations associated with the final results should be combined and enforced.

R31. In a distributed environment the obligations potentially associated with a response should be specified.

R32. The relation between obligations should be specified in order to support their combination.

While most of these requirements hold true for the non-distributed case, the combination of results from the authorization mechanisms is only necessary when the decision occurs in more than one place.

6.2.3 Enforcement

R33. Network access control policies have to be reliably enforced

When an End User is roaming, the visited institution or network is responsible for applying a particular resource access policy according to the agreements derived from the federation, therefore network policies should be applied in the network the End User is attached, even in roaming.

R34. The enforcement process of access control policies should support negotiation which aimed at establishing the least set of information that a user want and has to disclose before accessing a specific service.

6.3 Telecommunications Requirements

R35. All communication must be identity-bound.

The framework should be symmetric in its use of digital identities for all transactions, including communication between two or more peers. The digital identities should be available to all allowed network and service elements to use as references to information about users, services or, in fact, any object.

R36. Transactions among the Identity Management framework and users, service or network elements must provide authenticity, integrity and confidentiality.

Conversely, non-repudiation may also be provided in some cases.

R37. Bi-directional authentication of requesting authorities and Provisioning Service Points.

R38. Mutual authentication must be performed so a trust relation can be established.

In order for a trust relation between two Provisioning Service Points of different Provisioning Domains to exist those two provisioning points must perform mutual authentication so as to ensure that no identity related data is divulged to unknown third parties.

R39. Previous roaming agreements should exist between different operators.

A roaming agreement is part of the agreement between operators and allows them to exchange data and/or decision information. While in many providers this agreement can be made ad-hoc, for an operator it is many times important to establish the agreements *a priori*. When such an agreement is established, the boundaries for distributed access control should also be defined.

R40. Decision and enforcement points have to be clearly defined and functionally independent.

R41. Access control decision and enforcement functions may be present in different layers (transport, control, service).

Access control may occur in many levels of the service provisioning. In the case of an operator, policies may define generally access control rules which affect network and service enforcement points.

R42. The Access Control entities functionality and its distribution should not limit the inclusion of new business models.

The extensibility of the framework is crucial to support new services and business models. As such, access control language and mechanisms should remain at a level which allows for an easy integration with existing or new services.

6.4 Access Control and Identity Management Requirements

R43. As one component in the Identity Management lifecycle, the use of credentials (containers for identity information e.g. digital certificates) for identifying, authenticating and authorizing user for access to protectected objects and ressources has to be in compliance with its privacy preferences.

R44. Architecture must be scalable with particular attention to user centric Identity Management mechanisms.

The IdM framework should be designed with scalability in mind in terms of: bandwidth, real-time data, including low-latency requirements, redundancy and high volume of requests per second. The IdM framework must also be accessible from many different service and network elements. In addition, it should support user empowerment to manage their privacy: for instance by allowing them to specify enforcable data handling restrictions and constraints about the use of their personal information.

R45. The Identity Management framework must not disallow legacy services (non-framework enabled services).

The Identity Management architecture should evolve and adapt its protocols to new standards and practices while maintaining the interfaces to legacy services and applications.

R46. Unique and precise discovery of identity resources and attributes must be provided.

The architecture should enable simple, consistent and scalable ways of discovering and resolving information relating to the interacting identities, enhancing the communication process in several planes of operation.

R47. Identifiers should be dynamically generated.

The identifiers used across the architecture should be dynamically extractable from the identity namespace, and not impose restrictions on the resolution or the storage processes.

R48. Identifier generation should be privacy aware, but sill provide useful information

Identifier generation, allocation and usage should depend on the current identity and on volatile information, such as context personalization and, most importantly, privacy.

R49. The authentication context and authentication token shall support different methods of multi-factor authentication, including current, standardized authentication methods as well as future ones.

Several authentication methods must be able to be combined in order to increase quality or efficiency of the authentication. Established authentication methods must be supported for a migration process.

R50. Services must be securely separated for controlled delegation of access rights.

If a user authorizes anyone else to use a certain service on his behalf, this authorization must be automatically restricted to this service. This means that different VIDs with different delegation credentials must be presented for each service.

R51. The Identity Management architecture must ensure high availability. Furthermore, access control mechanisms should not be bypassed even when interacting entities wish, for privacy reasons, to limited or prevent the disclosure of personal information.

The operational site has to implement a fully working operating concept to ensure that no process can influence the availability of the services. The architecture must be distributed to withstand maintenance, security and deployment disruptions, ensuring that there is no single point of failure.

6.5 Summary Table of Requirements and Map to Use Cases

Table 2

Number	Summary	Map to Use Cases
6.1 General Access Control Framework Requirements		
R1	In order to support transparency, there should be a mechanism for an entity to provide evidence that it needs certain information from the user and an interface for external auditing in terms of privacy policies and data processing.	**UC1, UC2, UC3**
R2	Authentication, Integrity and non-repudiation should be enabled for all transactions.	**UC1, UC2, UC3**
R3	Support of granular authorization.	**UC1, UC2, UC3**
6.1.1 General Access Control Framework Requirements: Policy Management		
R4	Authenticity, integrity and non-repudiation should exist between the different entities.	**UC1, UC2, UC3**
R5	Users should have a simple mechanism to both set and realize the consequence of policies; even when these policies are set by an agent on behalf of the user.	**UC2**
R6	Enable authorized personnel to audit the status and usage of the security mechanisms.	**UC2, UC3**
R7	Availability of Preferences.	**UC2, UC3**
R8	The framework must support dynamic management of policies at any time.	**UC2, UC3**
R9	The access control framework must support the delegation of rights.	**UC2, UC3**
R10	The possession of attributes must be unforgeable.	**UC1, UC2, UC3**
6.1.2 General Access Control Framework Requirements: Decision		
R11	Authentication assertion and authentication context should be available for the authorization.	**UC1, UC2, UC3**
R12	If an authentication assertion could not be directly understood by the original requestor, a method to transform the assertion and related data should be provided by trusted entities.	**UC1, UC2, UC3**
R13	The requestor is authenticated and is either a user, an application acting on behalf of a user, or a machine running an application and/or under the control of a particular user.	**UC2, UC3**
R14	The authorization for a particular type of access should be based on a request which includes related attribute information and the resource with related information.	**UC1, UC2, UC3**
R15	Authorization requests should be responded within a well defined time frame, or a default reaction should be enforced.	
R16	Authorization responses may include addition obligations which have to be enforced as a reaction of the request independently whether the response was a denial or a permit.	**UC1, UC3**
6.1.3 General Access Control Framework Requirements: Enforcement		
R17	User agent should be able to authenticate to a mutually agreed authentication server.	**UC2, UC3**
R18	Different types of authentication technologies or protocols can be supported.	**UC1, UC2, UC3**
R19	Authentication request might be forwarded to another authentication server.	**UC3**
R20	An authentication server function should exist and should be able to create assertions about the user's identity.	**UC1, UC3**
R21	Consistent policy enforcement must be available on each layer of the architecture.	**UC1, UC2, UC3**
6.2 Distributed Access Control Requirements		
R22	Establishment of Trust Relationship.	**UC1, UC2, UC3**
R23	No spread of security breaches.	**UC1, UC2, UC3**
R24	Retrieval of attributes from several different Attribute Providers must be possible.	**UC1, UC2, UC3**
R25	The framework must support the combination of distributed or cascaded policies from different administrative entities.	**UC1, UC2, UC3**
6.2.1 Distributed Access Control Requirements: Policy Management		
R26	A central point collecting all the policies of different entities should be avoided.	**UC1, UC2, UC3**

Number	Summary	Map to Use Cases
R27	Identity management framework must provide the services and users the way to discover Identity Brokers (for Single Sign On/Single Log Out) and Attribute Providers (for attribute exchange), and obtain user's attributes from them under user's control, with using user's pseudonym or anonym.	**UC2, UC3**
R28	The policy-based access control framework should provide means for managing the overall policy life cycle, i.e. by providing functions for specifying, monitoring, enforcing and de/activating policies or providing mechanisms to guarantee the secrecy of policies (since sensitive information related to the policy can be deduced from the exchange between interacting entities even when the policy itself is not disclose).	**UC2, UC3**
6.2.2 Distributed Access Control Requirements: Decision		
R29	In a distributed environment authorization decisions may depend on decisions of other entities. The requesting entity is responsible for combining the results.	**UC1, UC2, UC3**
R30	In case the final decision depends on multiple decisions by different entities all the obligations associated with the final results should be combined and all obligations should be enforced.	**UC1, UC2, UC3**
R31	In a distributed environment the obligations potentially associated with a response should be specified.	**UC1, UC2, UC3**
R32	The relation of the obligation should be specified in order to support their combination.	**UC1, UC2, UC3**
6.2.2 Distributed Access Control Requirements: Decision		
R33	Network policies should be applied in the network.	**UC1, UC2, UC3**
R34	The enforcement process of access control policies should support negotiation which aimed at establishing the least set of information that a user want and has to disclose before accessing a specific service.	**UC1, UC2, UC3**
6.3 Telecommunications Requirements		
R35	All communication must be identity-bound.	**UC1, UC2, UC3**
R36	Transactions among the Identity Management framework and users, service or network elements must provide authenticity, integrity and confidentiality.	**UC2, UC3**
R37	Bi-directional authentication of requesting authorities and Provisioning Service Points.	**UC1, UC2, UC3**
R38	Mutual authentication must be performed before a trust relation is established.	**UC1, UC3**
R39	Previous roaming agreements should exist between different operators.	**UC3**
R40	Decision and enforcement points have to be clearly defined and functionally independent.	**UC1, UC2, UC3**
R41	Access control decision and enforcement functions may be present in different layers (transport, control, service).	**UC1, UC2, UC3**
R42	The Access Control entities functionality and its distribution should not limit the inclusion of new business models.	**UC1, UC3**
6.4 Access Control and Identity Management Requirements		
R43	As one component in the Identity Management lifecycle, the use of credentials (containers for identity information e.g. digital certificates) for identifying, authenticating and authorizing user for access to protected objects and resources has to be in compliance with its privacy preferences.	**UC1, UC2, UC3**
R44	Architecture must be scalable with particular attention to Identity Management user centric mechanisms.	**UC2, UC3**
R45	The Identity Management framework must not disallow legacy services (non-framework enabled services).	**UC1, UC2, UC3**
R46	Unique and precise discovery of identity resources and attributes must be provided.	**UC2, UC3**
R47	Identifiers should be dynamically generated.	
R48	Identifier generation should be privacy aware, but sill provide useful information.	
R49	The authentication context and authentication token shall support different methods of multi-factor authentication, including current, standardized authentication methods as well as future ones.	**UC1, UC2, UC3**
R50	Services must be securely separated for controlled delegation of access rights.	**UC1, UC2, UC3**
R51	The Identity Management architecture must ensure high availability.	**UC1, UC2, UC3**

7 Conclusion

While we have identified some of the major requirements in having a distributed access control which spans over applications, services and networks. The link between the access control framework and the identity management framework is evident by the close link between data, access, service and user.

In the next steps, having identified priorities in services influencing network policies and operations, we must reuse the protocols identified in clause 4, bridged with extensions, to fulfil the requirements identified in clause 6.

Annex A (informative): Bibliography

XML Document Management (XDM) Specification, Version 1.0, Open Mobile Alliance™, OMA-TS-XDM-CORE-V1-0.

NOTE See http://www.openmobilealliance.org/

"PoC XDM Specification", Version 1.0.1, Open Mobile Alliance-, OMA-TS-POC-XDMV1-0-1.

NOTE: See http://www.openmobilealliance.org.

History

Document history		
V1.1.1	September 2010	Publication