

Proposal for a Common Identity Framework:

A User-Centric Identity Metasystem

Kim **Cameron**
Reinhard **Posch**
Kai **Rannenber**g

Oct 05, 2008

1. TABLE OF CONTENTS

2.	Introduction.....	4
3.	Terminology.....	5
4.	Scope	7
5.	Metasystem Requirements in the light of Multilateral Security	8
5.1	Requirements of sites and services using the system.....	8
5.2	Requirements of people using the system.....	9
5.2.1	Strict Control of information flows by users.....	9
5.2.2	Data Minimization	9
5.2.3	Contextual Separation	9
5.3	Information protection framework.....	10
5.4	Freedom of choice	10
5.4.1	Choosing Operators	11
5.4.2	Interoperability	11
5.5	Requirements of governments	11
6.	Abstract Model of the Identity Metasystem	13
6.1	Claims.....	13
6.2	actors participating in the Metasystem	15
6.2.1	Subject	15
6.2.2	Claims Providers	15
6.2.3	Relying Party	15
6.2.4	Subject Acting As (SAA).....	15
6.2.5	Technical Policy Provider	15
6.2.6	Administrative Domain	15
6.3	Metasystem Agents and information flow.....	16
6.4	Contractual agreements between parties	17
6.5	The Abstract Services of the Metasystem.....	17

6.5.1	Primordial Claim Abstract Service.....	18
6.5.2	Registration Abstract Service.....	18
6.5.3	Claims Provider Abstract Service	19
6.5.4	Claims Selector Abstract Service.....	20
6.5.5	Claims Approver Abstract Service.....	21
6.5.6	Resource Matching Abstract Service	23
7.	Enabling Technologies	25
7.1	Minimal Disclosure Tokens	25
7.2	Minimum Footprint Technologies.....	26
8.	Administration	27
8.1	Administrative Domains.....	27
8.2	Definition of technical Policy.....	29
8.3	Enforcement of technical Policy.....	29
8.4	Auditing of technical Policy enforcement	29
9.	Standardization.....	30
9.1	What needs to be standardized?	30
9.1.1	Systems to support all the parties	30
9.1.2	Claims.....	30
9.2	Who should standardize what?	30
9.2.1	Properties of organizations.....	30
9.2.2	Sector specific vs. general.....	30

2. INTRODUCTION

This paper proposes a framework for protecting privacy and avoiding the unnecessary propagation of identity information while facilitating exchange of specific information needed by Internet systems to personalize and control access to services. It also sets out factors to be taken into consideration when deciding where the standardization of such a framework should be brought about.

Information systems that co-operate to originate, control and consume identity information have been called identity systems. The evolution of the Internet requires increased interoperability of these systems. Such interoperability demands an abstract model that encompasses the characteristics of all co-operating identity systems. We call this abstract model the Identity Metasystem.

Describing, designing, deploying and managing identity systems in accordance with this model will facilitate the interworking of identity components:

- from different manufacturers;
- under different managements;
- of different levels of complexity;
- based on different protocols ;
- employing different syntaxes;
- conveying different semantics; and
- of different ages.

3. TERMINOLOGY

The following concepts are employed:

- **Abstract services:** Architectural components that deliver useful services and can be described through high level goals, structures and behaviors. In practice, these abstract services are refined into concrete service definitions and instantiations.
- **Administrative authority:** An organization responsible for the management of an administrative domain.
- **Administrative domain:** A boundary for the management of all business and technical aspects related to:
 1. A claims provider;
 2. A relying party; or
 3. A relying party that serves as its own claims provider
- **Application Specific Identifier (ASID):** An identifier that is used in an application to link a specific subject to data in the application.
- **Claim:** an assertion made by one subject about itself or another subject that a relying party considers to be “in doubt” until it passes “Claims Approval”
- **Claims Approval:** The process of evaluating a set of claims associated with a security presentation to produce claims trusted in a specific environment so it can be used for automated decision making and/or mapped to an application specific identifier.
- **Claims Provider:** An individual, organization or service that:
 1. registers subjects and associates them with primordial claims, with the goal of subsequently exchanging their primordial claims for a set of substantive claims about the subject that can be presented at a relying party; or
 2. interprets one set of substantive claims and produces a second set (this specialization of a claims provider is called a claims transformer). A claims set produced by a claims provider is not a primordial claim.
- **Claims Selector:** A software component that gives the user control over the production and release of sets of claims issued by claims providers.
- **Claims Transformer:** A claims provider that produces one set of substantive claims from another set.
- **ID-data base:** A collection of application specific identifiers used with automatic claims approval
- **Identity:** The fact of being what a person or a thing is, and the characteristics determining this.
- **Natural person:** A human being
- **Person:** an entity recognized by the legal system. In the context of eID, a person who can be digitally identified.
- **Persona:** A character deliberately assumed by a natural person
- **Primordial Claim:** A proof – based on secret(s) and/or biometrics – that only a single subject is able to present to a specific claims provider for the purpose of being recognized and obtaining a set of substantive claims¹.
- **Registration:** The process through which a primordial claim is associated with a subject so that a claims provider can subsequently issue a set of claims about that subject.
- **Relying party:** An individual, organization or service that depends on claims issued by a claims provider about a subject to control access to and personalization of a service.

¹ The word primordial is used to refer to “first claim” in the sense of “constituting a beginning; giving origin to something derived or developed.” We have chosen to avoid the word “credential” in this regard given that it means many things, including both primordial and substantive claims.

- **Security presentation:** A set consisting of elements like knowledge of secrets, possession of security devices or aspects of administration which are associated with automated claims approval. These elements derive from technical policy and legal contracts of a chain of administrative domains.
- **Security Token:** A set of claims.
- **Service:** A digital entity comprising software, hardware and/or communications channels that interacts with subjects.
- **Subject:** The consumer of a digital service (a digital representation of a natural or juristic person, persona, group, organization, software service or device) described through claims.
- **Substantive claim:** A claim produced by a claims provider – as opposed to a primordial claim.
- **Technical Policy:** A set of technical parameters constraining the behavior of a digital service and limited to the present tense.
- **User:** a natural person who is represented by a subject.
- **User-centric:** Structured so as to allow users to conceptualize, enumerate and control their relationships with other parties, including the flow of information.

4. SCOPE

In addition to defining a model, the proposed framework defines abstract services facilitating interoperation of Identity Metasystem components. Such services can be instantiated and optimized through given protocols and semantics, but such considerations are outside the scope of this discussion.

The specific features required in managing identity and access within given administrative boundaries may differ. Due to these differing requirements and various historical reasons, identity systems with different properties exist. But all of these systems conform to various degrees with the Identity Metasystem model, and to this extent can be made to interoperate. Such systems, which will continue to evolve, are the “constituent systems” of the Identity Metasystem. The integration of constituent systems through their own protocols being a key issue, the proposed framework would also describe mechanisms for making this possible.

The content of the information flows in the Metasystem is constituted of semantic fields that are of interest to interoperating parties in government, industry and commerce, as well as to other stakeholders, and importantly, to the individuals about whom information is exchanged. The framework therefore includes a mechanism for mapping semantic fields as “claims”. Their content is open-ended, and the model is modular and flexible in that independent domain-specific initiatives can address these problems adaptively (e.g. in government, industry verticals, academia, etc).

Given the importance of personalization and access control to future digital life, the Metasystem framework can be expected to become the basis for many standards and recommendations involving identity information.

5. METASYSTEM REQUIREMENTS IN THE LIGHT OF MULTILATERAL SECURITY

Organizations that offer digital services and operate Internet web sites, as well as individual users, have numerous requirements with regards to the features and governance of the Identity Metasystem.

The prevention of online fraud and identity theft is a central goal. So is protection of the privacy of individuals and organizations.

The user-centric Identity Metasystem is underpinned by three important concepts: transparency, consent and security. These should be enforced through the use of technology to enable:

- a secure infrastructure: employing safeguards that help protect against malware and unauthorised access to personal information, and that help keep systems up-to-date;
- strong identity and access control: systems that help protect personal information from unauthorised access or use.

5.1 REQUIREMENTS OF SITES AND SERVICES USING THE SYSTEM

A major interest of the organizations that operate web sites and services is to ensure that users get access to personalized services and resources while unauthorized or fraudulent parties do not.

As the needs of business and organization change, and as partnerships and alliances evolve, it is necessary for new systems to be able to interwork without modification at the infrastructural level. Sites need to be able to adapt flexibly as their purposes and interests change.

In addition, it is a goal that risk and liability be reduced. One example would be for the sensitive information in an enterprise or government department to be quarantined rather than propagated throughout back office systems, reducing the probability of incurring damages should information leak or be abused. Another would be for an organization to avoid asking for and holding a subject's name, address, and national identifier: the Metasystem allows a relying party to substitute a "derived claim" – e.g. an assertion by a trusted party that the subject resides in a given city, is a citizen, and has a valid national identifier – without requiring the national identifier to be stored. This example demonstrates the advantages of depending on strong authentication rather than the propagation of sensitive information: this approach provides numerous benefits in terms of protection from identity theft, fraud and insider attacks and "data loss". The discussion of [Data Minimization](#) points to various ways the system can be structured to accomplish these purposes.

Finally, it is highly desirable that compliance with relevant statutes, standards and audit requirements be an automatic outcome of the Identity Metasystem as instances are deployed.

5.2 REQUIREMENTS OF PEOPLE USING THE SYSTEM

There are four main interests from the user side:

- 1) Co-operating to ensure resources associated with the user are protected from unauthorized access.
- 2) Being able to control and benefit from information flows
- 3) Enjoying data minimization
- 4) Achieving a separation of contexts on par with that characterizing the physical world.

All of the interests have strong relation to users' privacy.

5.2.1 STRICT CONTROL OF INFORMATION FLOWS BY USERS

The core requirement for user control is that the flow of information from Claims Providers to Relying Parties only happens at the request of the user. This has two major aspects:

- 1) Human Factoring: the presentation of human interfaces that are convenient and unambiguous.
- 2) Transparency and disclosure towards the users, who need at all times to understand and control what information is being exchanged and for what purpose.

5.2.2 DATA MINIMIZATION

Data Minimization applies to all the processes that deal with personal data. All of the following processes should work with the minimum amount of personal data and be designed in that way:

- Collection
- Aggregation
- Storage
- Retention
- Replication
- Distribution
- Linkage

5.2.3 CONTEXTUAL SEPARATION

The need for Contextual Separation is a corollary of Data Minimization, since the introduction of links between activities in different contexts is a form of aggregation and collection.

For example, Data Minimization implies that the relationships of consumers with different enterprises should not be amalgamated into super-dossiers. Nor should consumer information be integrated with government information. Indeed, activities with unrelated government departments should be kept separate. The concept of "Partial Identities" developed in the FIDIS and PRIME projects addresses this requirement, and "Partial Identities" can be modeled via the Identity Metasystem described in this text.

In particular, the use of the same identifier across different unrelated contexts is incompatible with the requirements defined in 5.2.2.

5.3 INFORMATION PROTECTION FRAMEWORK

There is a large body of work on Information Protection and Information Assurance. Here we are concerned with how the Identity Metasystem ties into and supports this work.

Beyond the critical decisions about what is to be collected and stored, there are architectural and technical mechanisms that should be brought to bear to achieve the goals of data protection.

Minimizing the risk of leaking claims about people and organizations is a fundamental privacy requirement of the digital world and an underlying principle of abstract Identity Metasystem design. Conforming to this principle protects users, relying parties and identity providers.

Concrete Identity Metasystem components should be designed on the basis that breaches will occur. During breach, systems must leak the minimum possible information. Threat modeling, risk analysis and demonstration of conformance with the principles outlined here should become standard parts of deployment practice.

The mechanisms of encryption, access control, separation of duties, auditing and physical control are all absolutely necessary when dealing with identity information.

In addition, four partitioning approaches are especially important to minimizing the impact of any breach:

- 1) Reducing the number of collocated records;
- 2) Reducing contents of each record;
- 3) Controlling access to these records based both on application and role;
- 4) Separation of identifiers that link directly to natural persons from other information.

Aggregation should only be done in light of specific needs and under strict control. Aggregated data collections, if they exist, should only be accessed by systems with a demonstrable requirement, and persist only as long as necessary

Audit information should be collected in encrypted form and otherwise protected such that it is only available to system components with demonstrable need to access it, as well as, to the extent possible, the subjects to which it pertains.

This Information Protection Framework could be formalized so as to provide an anchor for service and system providers to claim compliance (similar to ISO 9000), e.g. by publishing where they position themselves within the framework. Business partners, government and consumers could take this into account when deciding who to deal with.

5.4 FREEDOM OF CHOICE

Freedom of choice for both users and relying parties refers to choice of service operators they may wish to use as well as to the interoperability of the respective systems.

5.4.1 CHOOSING OPERATORS

Users need the freedom to choose operators from a number of context-specific operators as well as more general operators.

5.4.2 INTEROPERABILITY

Interoperability is a prerequisite for choice. It allows the use of multiple technologies as well as the use of multiple platforms and devices from multiple vendors while shielding users and system programmers from having to understand the underlying differences. In particular, the framework services at 6.3 and 6.5 aim at enabling choice through interoperability.

5.5 REQUIREMENTS OF GOVERNMENTS

Governments have unique requirements and responsibility when it comes to the identities of natural persons. In democratic countries, citizens have established their governments and have asked them to make, interpret, and enforce law and policy. In this arrangement, governments control resources of great sensitivity and unquantifiable value. Hence, digital identity must be understood in this historical context: Governments have had control over resources before the existence of digital relationships, and some branches of the state have had unequaled access to personal and behavioral information.

All this implies the need for many rigorous controls, yet the requisite architectural components are the same ones needed in private enterprises.

[Claims Approval](#) and [Resource Matching](#) may be especially stringent given the fact that there is no obvious way to compensate for damage that might accrue from errors in this regard.

Primordial Claims may be associated, for example, with governmental identity cards, to help provide reliability in protecting citizens' resources from those who should not have access to them.

Similarly, registration may involve in-person proofing, and even require periodic renewal.

Yet these strong registration processes and primordial claims mechanisms make it possible to eliminate the release of personally identifying information - including linkable identifiers – when gaining admission to many services. This is explained in the section on Enabling Technologies.

The data protection and minimization precautions necessary in any identity system apply even more strongly to government systems, given the sensitive nature of the information and the difficulty of adequately compensating its compromise.

There may be a single or multiple government claims providers operating on behalf of different levels of government and associated with different departments (e.g. Health versus Travel). These may be represented through multiple digital cards within a claims selector and produce unlinkable claims.

A complication arises from the fact that government may, in some cases, appoint proxies to act on behalf of citizens (for example, on behalf of citizens who are mentally infirm, disadvantaged in terms of technology access, imprisoned and the like).

Thus it may be necessary for some information about Application Specific Identifiers and resource content to be available to specialized agents of government within constitutional limits.

The essence here is to preserve the normal data minimization and cross-context separation aspects described in the [Information protection framework](#).

For example, the fact that some agents need access to information must not mean the information is generally available and data minimization requirements do not apply. It should only dictate that specialized agents may be legitimate parties to protected information within some constrained scope.

6. ABSTRACT MODEL OF THE IDENTITY METASYSTEM

In light of these requirements, the Identity Metasystem model defines:

1. A mechanism, called claims, for describing subjects, that works across all constituent identity systems;
2. A taxonomy of claims;
3. A taxonomy of parties present in the system, including subjects,
4. The components through which the users interact with the system;
5. The abstract services offered by the components
6. The privacy and security threats arising from the information flows.
7. The system requirements arising from these threats
8. The establishment and use of technical policies

It also calls attention to the need for a complementary legal framework.

6.1 CLAIMS

A claim is an assertion made by one subject about another subject that is defined to be “in doubt” until passing “Claims Approval”.

By doubt we mean:

1. The integrity and origin of the claim needs to be verified (e.g. through cryptography and evaluation of a security presentation); and
2. The meaningfulness of a given party making a given claim about a given subject needs to be determined.

Through cryptographic methods, “doubt” may be resolved without any need to “call home” to the subject’s claims provider.

Subjects may be individual people as they exist in various contexts, groups, organizations, enterprises, governments, agencies, digital services and devices.

The degree to which a relying party is willing to believe or act upon a claim from an originating party constitutes part of a relying party’s technical policy. Elaboration of this technical policy is the responsibility of the relying party’s administrative domain (see 6.2).

The taxonomy of claims includes:

Type of Claim	Comment	Examples
Static	What we have traditionally called “properties” and “attributes” of the subject – static within <i>some window</i> of time	National identifiers and employee numbers Date of Birth Name Address
Relationship	Subject is in some relationship with another subject (and open-ended model with multiple sources and viewpoints)	Member of arbitrary group Member of assigned role Relationship to another subject (e.g. Personal Assistant or Parent) Mandate (e.g. trustee) Acting-as / On-behalf-of relationships
Derived	Claims that convey minimum necessary information by deriving it from facts but not releasing the facts	Over 21 or Under 16 University Student Person in Drug Trial Unmarried Female in 20’s
Capability	Authentication and authorization both based on claims transformation. Capabilities are determined by relying party within a defined scope	Can-read-calendar Can-access-write-operation Denied-update-in-given-scope
Contextual Claims	Factors useful in evaluating the security presentation.	Authentication technology, location, time

Constituent identity systems can all be reduced to systems for conveying claims. In particular:

- Kerberos² and similar protocols convey the claim that a subject has a given identifier within some domain (and possibly related attributes such as group membership made possible through an extension mechanism);
- Public Key³ Infrastructures transfer claims about the names and keys of subjects, as well as other identifiers and an extensible set of attributes;
- SAML⁴ conveys assertions which are a set of claims;
- OpenID⁵ uses the DNS infrastructure to validate the claim linking a subject to a URL.

² RFC 1510 - The Kerberos Network Authentication Service (V5)

³ RFC 3280 - Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List Profile

⁴ See OASIS Security Services (SAML) TC

6.2 ACTORS PARTICIPATING IN THE METASYSTEM

The actors participating in the Identity Metasystem can be classified by role, taking into consideration that any individual actor or set of actors can play multiple roles (both at the same time and at different times).

6.2.1 SUBJECT

A subject is a consumer of a digital service. Subjects may act on their own behalf (as individual citizens, consumers or cyber dwellers), or in roles within organizations, enterprises or government departments. Devices and digital services are subjects acting on behalf of other subjects.

6.2.2 CLAIMS PROVIDERS

A claims provider is a digital service through which an individual or organization makes a claim about another individual, organization, device or service.

6.2.3 RELYING PARTY

A relying party is an individual, organization or service that depends on claims issued by a claims provider about a subject to control access to and personalization of a service.

6.2.4 SUBJECT ACTING AS (SAA)

An SAA is a subject that acts on behalf of another subject. One example would be a person who is given a “power of attorney” by another person. Similarly, government officials sometimes act on behalf of specific citizens. Another common case is that of digital services that act on behalf of other subjects.

6.2.5 TECHNICAL POLICY PROVIDER

A technical policy provider is an individual or organization that creates policies employed by a relying party (and its agents) to decide how claims should be translated into service permissions and personalization.

6.2.6 ADMINISTRATIVE DOMAIN

An Administrative Domain is an entity which operates and manages some set of Metasystem components, and is responsible for the functioning of those components, and for the development of legal contracts and technical policies governing the use of those components.

⁵ See OpenID Authentication 2.0 at <http://openid.net>

6.3 METASYSTEM AGENTS AND INFORMATION FLOW

Human users, including organizations, act in the digital realm through agents that operate on their behalf.

The parties to the Identity Metasystem described in 6.2 operate through agents as represented in Figure 1.

- 1) The user employs a computer agent (for example a web browser or software program) to consume services from a service provider (for example a web site or web service).
- 2) In response to a service request, the service provider may inform the user's agent, through a technical Policy requirement, that to grant access or personalize behavior it requires identity information about the user.
- 3) The user's agent presents that information to the user through a specialized agent called a claims selector. Should the user instruct the claims selector to release the required claims, it contacts a claims provider, conveying the relying party's technical Policy requirement. The claims selector also conveys a pre-arranged proof that it is operating on behalf of the user.
- 4) The claims provider uses its technical Policy to determine what claims it should issue (if any) given the proof supplied by the user's claims selector and the technical Policy requirement originating from the relying party. Resulting claims are returned to the claims selector.
- 5) The claims selector forwards the claims to the relying party. The relying party then uses its technical Policy to determine whether to recognize the user as a subject and how to personalize the subject's service. It may employ other agents to help make these decisions.

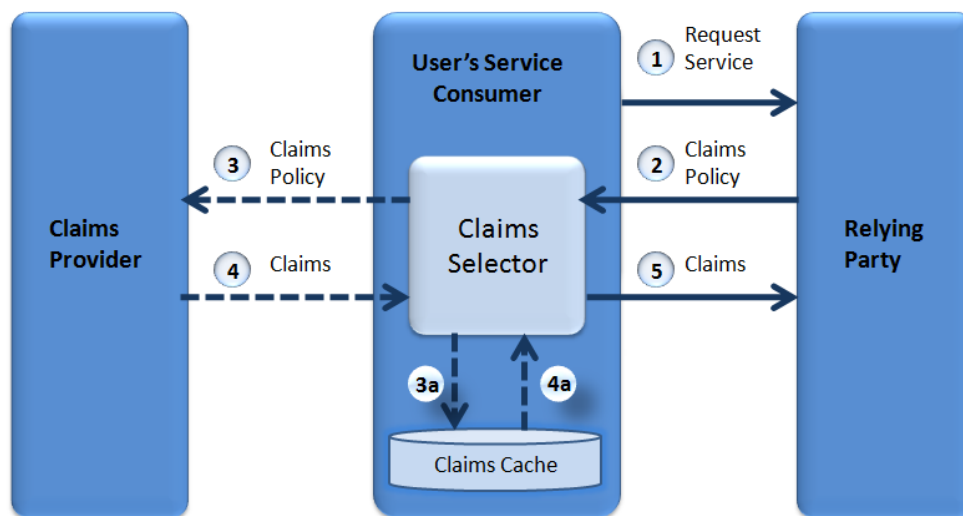


Figure 1 - Information flow in Identity Metasystem

This section defines an underlying pattern, and implementations may optimize the data flows. For example, there is no intent to constrain the lifetimes of sets of claims, which might be cached within the user's service consumer (3a), and once approved, subsequently presented to the relying party in an automated fashion (4a). In one variant of this, a claims selector acquires "packages of claims" from a claims provider in advance, along with a means of proving they pertain to a given user. These approaches potentially improve the performance, reliability and privacy characteristics of the system, as described in the section on Enabling Technologies.

6.4 CONTRACTUAL AGREEMENTS BETWEEN PARTIES

For the system to function effectively there can and sometimes must be explicit or implied contractual agreements between the parties. When the administrative authority operating a relying party decides to accept claims, it may do so under a contractual agreement with the administrative authority operating the relevant claims provider.

Amongst other things, such a contract would define:

1. usage restrictions and permissions
2. information quality
3. information protection assurances
4. auditing requirements
5. data minimization mechanisms as discussed in [Data Minimization](#)
6. quality of service
7. liabilities incurred
8. fee structure, etc.

However a relying party may decide to rely upon a claim even if no contractual relation with the claims provider exists, provided the identity provider is willing to issue it for such a use.

When both claims provider and relying party are operated within one administrative domain, these agreements become an internal matter.

There are similar contractual aspects to the relation between a subject and a claims provider that are agreed upon during registration, and between a subject and a relying party agreed upon when establishing or modifying a relationship with that entity. A legal and policy framework is required that will simplify the establishment of these agreements and concerted work by policy and legal experts needs to go into elaborating this framework.

6.5 THE ABSTRACT SERVICES OF THE METASYSTEM

The Metasystem can be factored into architectural components that deliver useful services and can be described through high level goals, structures and behaviors. We call these components abstract services, meaning they can be turned into concrete instances through “refinements” producing two broad outcomes:

1. protocols, syntaxes and ultimately software;
2. social and organizational mechanisms for service provision and consumption

The Identity Metasystem encompasses both authentication and authorization. However it distinguishes between two kinds of authentication:

1. the use of “primordial claims”, typically keys, to authenticate to a claims provider for the purpose of obtaining a set of claims made about the subject; and
2. the use of a set of claims about the subject to authenticate to digital services.

A claims provider may use a set of claims as an input to an authorization decision. It can then return the decision in another set of (authorization) claims.

6.5.1 PRIMORDIAL CLAIM ABSTRACT SERVICE

The Primordial Claim Abstract Service is the service through which a user (or service) generates a “primordial claim” that is the first input to the set of claims providers who produce the claims describing a subject.

A primordial claim can be employed solely by one subject. It can be thought of as a secret such as a password or key, but in practice systems employ a “function” of the secret - a digest or a signature – since this is more resistant to attacks.

The essence is that the primordial claim is not believed because it is asserted by some claims provider. It is accepted by a specific claims provider because, in a prearranged registration / provisioning process, the claims provider has ensured that a given digital subject is uniquely capable of employing it.

In this sense, typical smartcards including eIDs, one-time-password devices, trusted platform modules and even password entry subsystems all provide the Primordial Claim Abstract Service. They each use a secret known only to a given device or user.

In sufficiently controlled environments, some biometrics could also serve as inputs generating primordial claims (i.e. satisfy the condition of being able to be generated solely by one subject).

Primordial claims can be combined, as happens in the case of a smart card that requires a PIN. In this case, the PIN is a primordial claim used to access the card, which produces a second primordial claim in the form of a signature destined to a claims provider. In some cases contextual claims (e.g. location) may also be combined with primordial claims.

All mechanisms for generating Primordial Claims are vulnerable to attack. The mechanisms can be arranged on a spectrum ranging from the most vulnerable (e.g. user name and password) to the least (currently, tamper-resistant smart cards with biometrics and PINs). The security presentation is in part determined by where a subject’s primordial claims mechanism falls on this spectrum.

6.5.2 REGISTRATION ABSTRACT SERVICE

Registration is the process through which a Primordial Claim is associated with a subject so that a claims Provider can subsequently issue a set of claims about that subject. This process can be more or less stringent depending on the requirements of different contexts.

At one end of the spectrum, some claims providers might demand physical identification such as a birth certificate, driver’s license with photograph, banking information and passport or government identity card / social security number and background check to establish what claims can be made about a subject. In the registration process, this set of claims may then be associated with a primordial claim such as a key in a smart card. Subsequently, the data set, or some derivative, becomes the basis for a claims provider issuing claims (as described in Section 6.5.3) when the smart card is exercised.

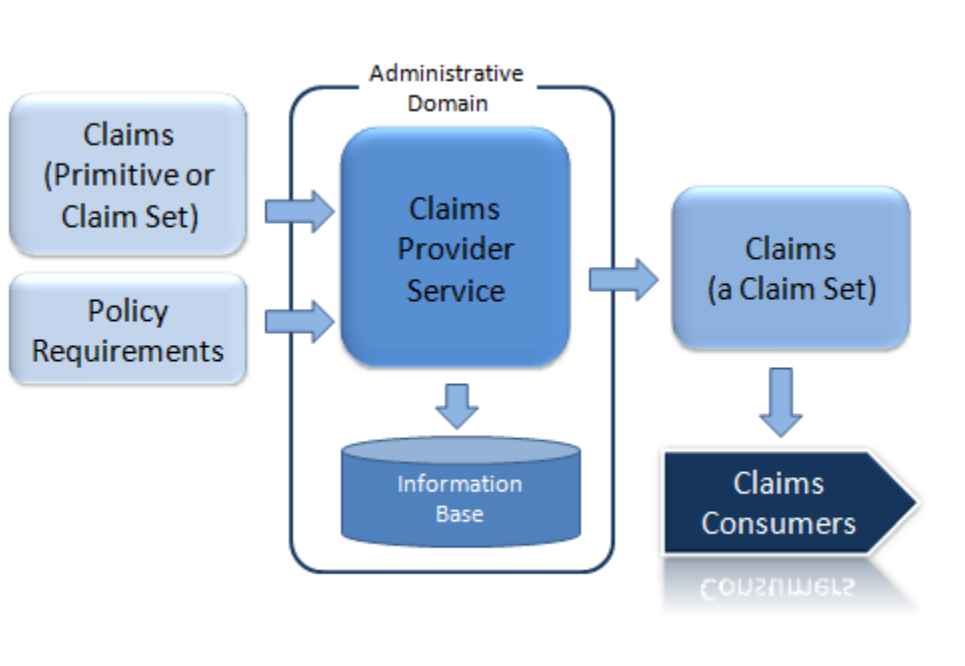
At the other end of the spectrum, registration can involve nothing more than the creation of secret password. The subject’s knowledge of this primordial claim can be used by the claims provider to link the subject to some (potentially pseudonymous) identifier or other substantive claimset..

Registration may also be an incremental process, beginning as pseudonymous and accruing identity information as appropriate to different circumstances in which the subject employs the claims provider.

Finally, registration with one claims provider can be bootstrapped through the claims issued by another. For example, the claims issued by a government claims provider could be used to register an employee with an enterprise claims provider. Subsequently the subject could employ a primordial claim provisioned by the enterprise claims provider to obtain enterprise claims, and the two digital identities would be independent going forward. The goal here is one of minimal disclosure, in which one set of claims is used to establish the second, and then, for purposes of contextual separation described in section 5.2.3, the relationship between the two digital identities is suppressed.

6.5.3 CLAIMS PROVIDER ABSTRACT SERVICE

The Claims Provider Abstract Service accepts one set of claims, along with a description of what claims are required, and issues a new set of claims.



The service can be thought of as a “claims transformation” service in which there are four inputs and one output:

Inputs	Claims	Primordial or Issued Claims provided with a service request
	Requirements	The technical Policy of a relying party or claims provider indicating information usage policy, what claims are required, the mechanism for expressing them, and other metadata that helps components rendezvous.
	Subject Information Base	Information maintained by the claims provider about digital subjects. The information may include the nature of relationships between subjects where a subject acts on behalf of another subject for specific services.
	Provider Technical Policy	A set of rules the provider employs to determine which claims are issued as a product of given input claims, requirements and facts.
Output	Issued Claims	A new set of claims (of any kind except primordial) which may potentially be the input to another claims provider

In actual refinements of the model, “optimizations” are possible. For example, Requirements or Technical Policy can be pre-arranged and hard wired. This is the case for most currently deployed identity systems. However, hard-wiring results in system isolation and the Metasystem framework is intended to move beyond this.

6.5.4 CLAIMS SELECTOR ABSTRACT SERVICE

The purpose of the Claims Selector Abstract Service is to give the user control over the production and release of sets of claims issued by his or her claims providers. Applications (relying parties) use the service to request such claims, along with any proofs necessary to show that it is entitled to submit them.

Inputs	RP Claims	Issued Claims identifying the Relying Party
	RP Requirements	The technical Policy the relying party indicating information usage policy, what claims are required, the mechanism for expressing them, and other metadata.
	Claims Selector Technical Policy Base	Information stored by the claims selector about user technical Policy decisions and preferences
	Relying Party Use Technical Policy	Statements on how the relying party will use any claims provided (privacy statement, information protection technical Policy, etc)
Output	Set of Claims (also called Security Token)	A set of claims of any kind which will be the input to a relying party or secondary claims provider

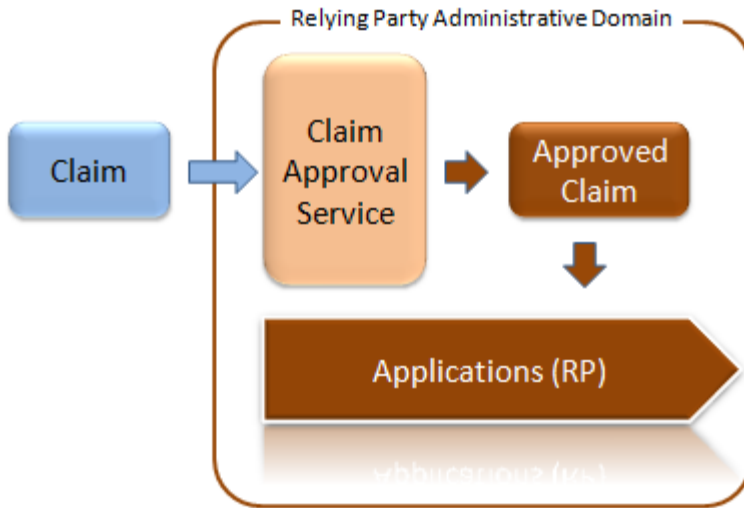
Claims selectors may be hard-wired as a result of policies agreed to outside the scope of the Identity Metasystem (most existing identity systems behave this way, and it has generally sufficed as long as interactions were taking place within a single administrative domain). However, as was the case with the Claims Provider abstract service, hard-wiring results in limited use patterns, isolated systems and loss of control and understanding by the user. The Metasystem model is intended to move beyond this.

The Claims Selector Abstract Service is the point of interface of the Identity Metasystem with its users. Service Requirements section of his document constrains the characteristics and operation of this service.

6.5.5 CLAIMS APPROVER ABSTRACT SERVICE

As claims are defined to be “in doubt” they are not to be relied upon until the relying party has decided to do so. This is called claims approval and results in a claim being transformed into an approved claim.

Factors potentially determining whether approval is given include purpose, technical Policy, digital integrity, security presentation, claim origination and subject, content, location and timeliness of the claims.



Claims approval is done by or on behalf of a relying party. A common scenario is one in which multiple services grouped within an administrative boundary depend on a single Approver service established to act on behalf of such sets of relying parties.

Inputs	Claims Content	Issued Claims identifying the Subject and the Originator of the claims.
	Approval purpose	Technical Policy statement of purpose for which claims are required
	Technical Policy	Factors constraining approval
	Security Presentation	E.g. - type of authentication and registration
	Claim metadata	E.g. - age of claims
Output	Approved Claims	A set of claims upon which a relying party can act

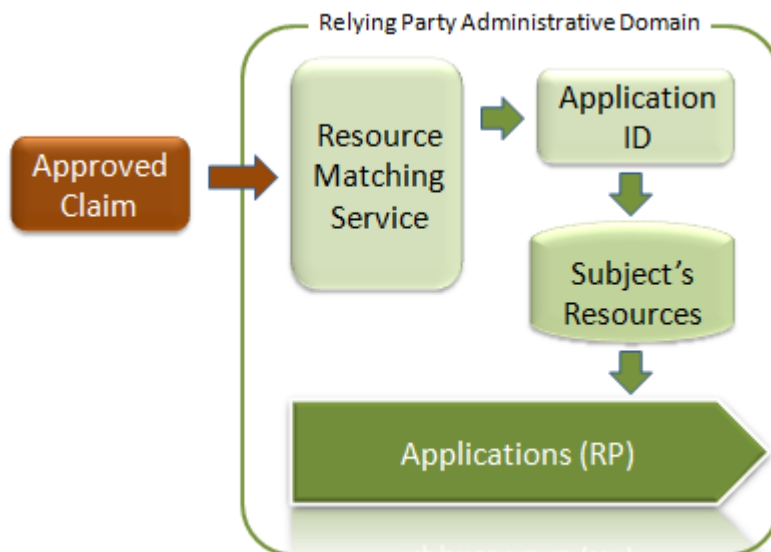
Applications can use the content of approved claims directly to shape the experience of their subjects. Examples might include selecting the language an application is presented in based on a claim about language preference; or configuring menu options based on a claim about a subject's roles; or controlling access based on employment details or age.

6.5.6 RESOURCE MATCHING ABSTRACT SERVICE

Some applications provide access to resources uniquely tied to the identity of a subject. For example, an on-line store might maintain information about each customer comprising purchase history, shipping information, “shopping cart” and wish list, and general preferences and interests. More dramatically, it is the responsibility of governments to ensure a tight binding between a digital subject and certain citizen entitlements and registries, e.g. for voting.

To make this possible, the Resource Matching Service connects one or more approved claims to the relevant subject resources. This is done by transforming an approved claim or set of claims to a local application identifier that serves to locate the subject resources within the boundary of an application.

The Resource Matching Service can be seen as a specialization and extension of a Claims Transformer – related but not identical even though the production of an identifier is involved.



That is because the Resource Matching Service functions in two modes: binding and access.

The binding mode involves the initial generation of an Application ID and connection of that ID to a set of resources. In some cases (for example, a new customer relationship) the set of resources might initially be empty – a “tabula rasa” – and this is a trivial exercise. In other cases there may be a valuable existing relationship between the application and the subject (for example, a land registry, health entitlement or pre-existing customer relationship) demanding strong verification and protection. In this case the binding mode acts to ensure that the right natural person is connected to the right set of resources. The binding mode may require presentation of a set of bootstrap claims sufficient to establish this mapping. For example, when a bidder or seller is setting up an eTrading account with a “high reputation”, the process could be streamlined and strengthened by submitting several claims from trusted authorities. “

Thereafter, the mapping can be represented through an Application Specific ID (ASID) bound to a potentially different approved claim.

In access mode, the ASID has already been established and connected to a relevant approved claim. Thus access to the application is done through a simple look-up of claim-to-ID. This is a streamlined operation that involves no exposure of personal information.

7. ENABLING TECHNOLOGIES

Technology and design will be crucial to meet the goals especially when it comes to security and data protection. If not designed in light of the kinds of technical and policy considerations outlined in this paper, an interoperable identity Metasystem would be more likely to erode privacy than to protect it; more likely to increase the problems of theft, fraud, insider abuse and coercion based on identity. The predictable result would be to reduce the public's confidence in the digital infrastructure. There is already widespread concern about the privacy implications of eID.

The architecture described here, combined with innovative uses of cryptography, provides an objective basis for dispelling these concerns. Beyond that, it promises significant reduction in the release of personally identifying information as compared with the status quo, and mechanisms to halt unnecessary comingling of profile data with claims identifying natural persons.

7.1 MINIMAL DISCLOSURE TOKENS

The principle of minimal disclosure ties information disclosed to what is demonstrably necessary for a transaction to complete. There are certainly situations in which significant information about a natural person is necessary for a transaction to be possible. This is the case, for example, when registering a deed.

But in general, our current systems overcompensate for the low quality of information and its uncertainty by collecting more information than is required. A “need to know” approach to transactions can only emerge based on confidence that things that are claimed can be counted on to be true.

Suppose the following conditions were met:

1. Existence of a set of organizations willing to make claims about subjects (for example, financial and governmental organizations).
2. The organizations employed high quality registration mechanisms resulting in a high degree of certainty about which natural persons they served.
3. The organizations were able to take advantage of strongly protected devices issued to users – whether in the form of smart cards or advanced embedded devices including phones.
4. Digital era financial and governmental services accepted claims issued by these organizations.

The architecture proposed allows users to contact their claim provider, authenticate through their strong device, pick up some “packages of claims”, and use their protected device to store them along with whatever proof is required to use them.

For example, one package of claims might allow a Belgian to pick up some claims saying she was a citizen, lived in Brussels, and was 32 years of age, along with a way of proving it.

When requesting services from a site that only serves Belgian citizens resident in Brussels, she would use her identity selector to present those claims, while suppressing the claim about her age since there is no “need to know” it.

A new cryptographic technology called Minimal Disclosure Tokens⁶ allows packages of claims to be created by the claims provider in such a way that:

1. they describe the registered subject to which they are given;
2. their authenticity and integrity cannot be tampered with;
3. they leak no information allowing the claims provider to track the usage of the claims - unless they are abused

It is possible to create strong disincentives to “lending” one’s claims to others.

Further, these “packages of claims” can be revoked if the claims provider has cause to do so. Yet users can demonstrate their claims have NOT been revoked without divulging ANY personally identifying information.

To see what could ultimately be achieved, one could create, for example, a digital passport that simply proved one was a Belgian Citizen AND not on a control list. The authentication would be much stronger than can be provided by today’s passports, and release no unnecessary information.

It should be clear that this solves many problems of today’s eID by introducing new privacy features that simultaneously increase the Multilateral Security of the system. At the same time, the approach lends itself to the wider Metasystem model because claims are no longer limited to particular hardware devices or national systems. Indeed, trans-border claims transformers can be put in place, and assuming the many problems of differing security presentations can be navigated, great progress can be made on international interoperability.

7.2 MINIMUM FOOTPRINT TECHNOLOGIES

Looking at a second example, people will want to take advantage of eID technology in a range of environments and in some cases, as users, will have limited control over the environment in use.

Minimum footprint technologies, that enable eID to be sandboxed elements virtualized and thus portable, offer the possibility of increasing the user’s confidence since the perimeter of use is clearly limited in time and in application domain spans. It is obvious that such technology has also to cover intermediate nodes and could be enhanced by combination with methods of the previous example.

Minimum or zero footprint technologies offer also to insulate domains from having to be aware of specific eID technologies as long as the tokens used talk the appropriate protocols.

A practical example would be a Spanish company starting an administrative process with an Austrian administration. A representative of that company could use its eID card and standard signature elements could be addressed by a virtualized security environment that temporarily downloads and encapsulates the eID function. Using SAML tokens as well as standard certificates usage of the Spanish eID including attributes becomes possible even cross borders without further registration at the Austrian eGovernment application.

⁶ Based on Zero-Knowledge proofs such as developed by David Chaum, Stefan Brands and Jan Camenisch and prototyped in PRIME and U-Prove

8. ADMINISTRATION

8.1 ADMINISTRATIVE DOMAINS

An Identity Metasystem administrative domain is the boundary for the management, deployment and operation of claims providers and/or relying parties. An administrative authority is responsible for the management of an administrative domain.

More specifically, an administrative authority is responsible for defining and managing Metasystem contracts, policies and operations, according to the model defined in this document, including the operation of standardized implementations of the relevant abstract services.

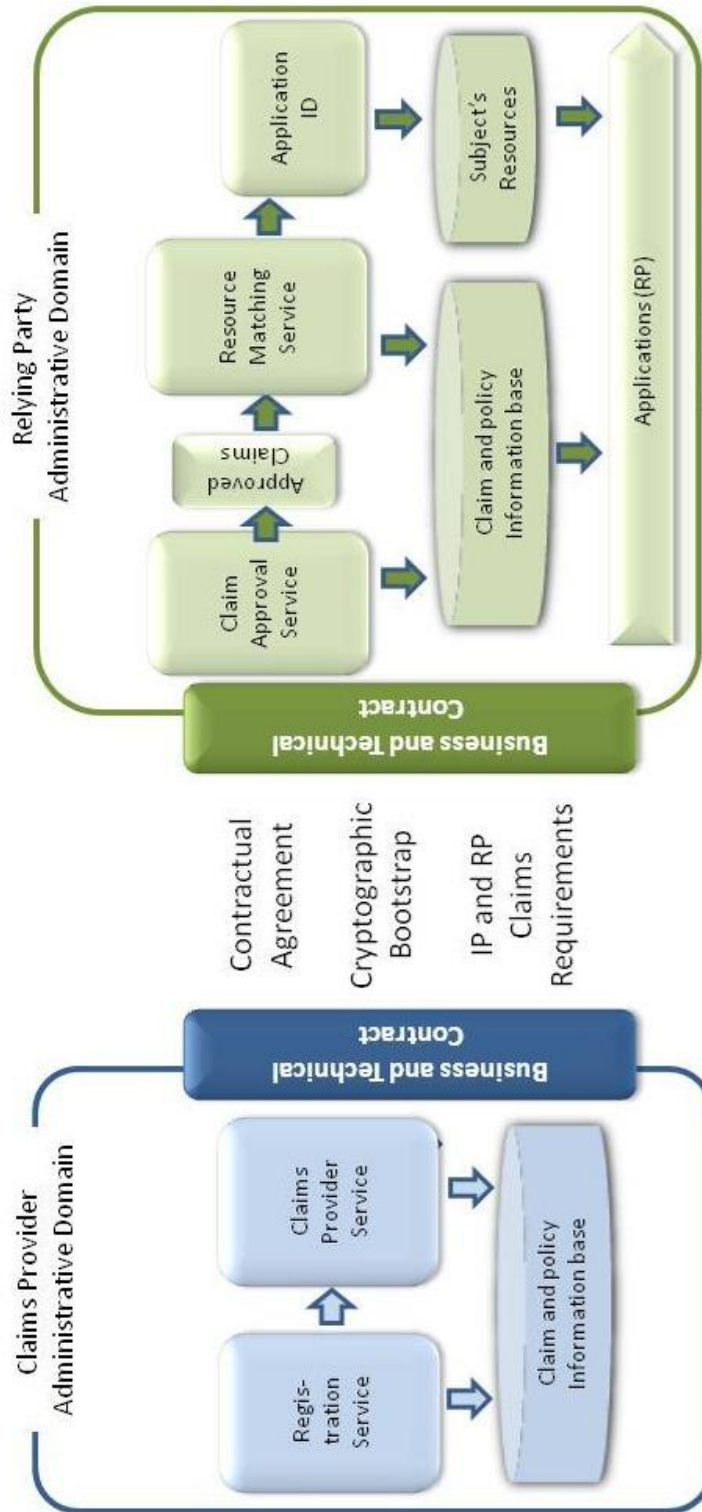
It is responsible for the quality of multilateral security, as delivered through the mechanisms defined in “Service Requirements in the light of Multilateral Security”.

A relying party and an identity provider may both live within a single administrative domain. One example would be an Internet service provider who offers access to a set of services through a portal and registers its own customers for portal access. Another example would be a relying party that operates a claims provider to transform external claims into a local format.

However, a relying party and a claims provider may also be located in different administrative domains. For example, this would be the case for a web site that federates with multiple enterprises; for enterprises that share resources across administrative boundaries; or for internet properties that accept claims made by credit card or financial providers.

When two or more administrative domains are involved, there is the requirement that the relying party and claims provider agree on business matters (e.g. nature of claims, quality of service, liability, business model, etc). There are also technical matters that must be negotiated and expressed through simple technical Policy.

Administrative domains can engage in multiple relationships governed by overlapping or distinct technical policies and legal contracts.



Note: This diagram does not show the partitioning required of information stores

8.2 DEFINITION OF TECHNICAL POLICY

By technical policy we mean a set of technical parameters constraining the behavior of a digital service and limited to the present tense. Examples would include the protocols understood, the claims required, and the information protection provided.

A key criterion for evaluating the success of policy statements is the extent to which they embody data minimization – for example derived claims.

Ensuring the use of derived claims and the structuring of policy to achieve data minimization needs to become a fundamental responsibility of officers of administrative domains, and be subject to audit and compliance requirements.

8.3 ENFORCEMENT OF TECHNICAL POLICY

Enforcement of technical policy requires:

1. Confidentiality mechanisms
2. Operational guidelines for identity information
3. A determination service to consume the claims and technical policy and provide decision and audit outputs

8.4 AUDITING OF TECHNICAL POLICY ENFORCEMENT

An auditing regime is required to verify the integrity of systems and data in compliance with business policies, and to ensure an organization can determine who has accessed a resource and who could access a resource.

9. STANDARDIZATION

9.1 WHAT NEEDS TO BE STANDARDIZED?

9.1.1 SYSTEMS TO SUPPORT ALL THE PARTIES

The framework components to be standardized are described in Chapter 6, and especially in 6.3 and 6.5.

9.1.2 CLAIMS

When claims are to be standardized this refers to Identity formats for claims and the packaging of claims

9.2 WHO SHOULD STANDARDIZE WHAT?

9.2.1 PROPERTIES OF ORGANIZATIONS

Four properties are essential to consider when deciding on the organization to standardize the Identity Metasystem:

- The ownership of the standards must be clear, and it must be with a respected and open organization.
- There must be a transparent, agreed and accessible process how to develop the standards and how to process amendments, e.g. in regular time intervals.
- There must be assurance that the "owning" organization lives long enough to not leave the standards as orphans.
- The standards must be neutral with regard to specific implementations and may also have to cope with regional and cultural differences (e.g. via allowing those as options).

9.2.2 SECTOR SPECIFIC VS. GENERAL

In principle standardization could be sector specific or general. Given the sector-overarching function of identity management it seems advisable to aim for "general" standardization instead of sector specific standardization. A typical example for general standardization is the Working Group (WG) 5 "Identity Management and Privacy Technologies" in Subcommittee (SC) 27 "Security Techniques" in the Joint Technical Committee (JTC) 1 "Information technology" of ISO and IEC.