

Creating a Relationship Binding

Abstract

Consider how humans create relationships; they first come together in some forum, physical or digital, identify the time and place of first meeting and start to exchange messages with language and with actions. This paper described how that process can work in an age of Intelligent Agents with the definition of a message to carry initial consent and a method for notification from a subject to any digital Service Provider with the outcome being an identifier for a shared context among a person and a collection of services and agents marshalled to solve a problem..

Goals

- Protect the data, rights and privileges delegated by a subject (the context) from reuse for purposes other than that intended by the subject with an identifier bound to that intent.
- The high-level goal is the establishment and maintenance of a relationship (binding) between the subject and a set of agents that can act on your behalf with the intent expressed and time-limited.
- It may be possible to add other subjects into a relationship, but there will always be (at least) one subject that “owns” the relationship and can shut it down, aka revoke all grants, causing the data held by the relationship to be destroyed.
- The immediate goal is for the subject to receive a consent receipt and authorization in response to this message which meets the security requirements of the intended purpose.
- Ensure solid binding to a secured user's key by requiring the user to have private components of the key well secured before this message is created.

You do have at least one account with your governments at all appropriate levels. It is an open question if a single government account is an acceptable privacy risk as they are just another leaking database.

- For healthcare the intended purpose is to bind the subject to the identity proofing document (IAL2) and validate the user agent as in compliance with any health care regulation or code of conduct (AAL2).
- The long-term goal is that no service provider would be permitted to retain any user information without first receiving a signed consent form like that described here and then remove all of that information when the relationship times-out or is cancelled.
- Allow a series of prompts that use results from prior prompts which also permits the subject to restart with a new relationship that has no consideration of any prompts in some other relationship.

Scope

- This specification applies only to service providers that have determined that they need to acquire explicit consent from individuals to create a shared information space.
- While a trust graph focuses on identification of the nodes, or individuals, a knowledge graph identifies the edges as well which correspond to the relationship between the nodes.
- The ceremony is the first interaction that creates user context within the larger user experience of Relationship Management.
- The term service provider is used here as the [Web Site](#) or Agent.
- It is possible for a subject to have a long term relationship with a limited set of user data and is augmented by short term relationships with more privileged but strictly limited time durations.

Pseudonomization

As noted there are many sources of an identifier from existing providers, typical email or mobile phone providers. Such a long-lived identifier must exist for functions like subscribing to any continuing service where the fee is fixed in some way that is independent of the service delivery. However the identifier is selected, just the continued reuse of that identifier will leak identity information.

In Trust-over-IP graph terms creating the relationship is described as creating a verifiable edge before trying to establish any identity for the nodes. See the paper "[Digital Contracts Made Legally Valid](#)" for a description on giving the relationship legal status.

Note that there are many names that are synonymous with relationship including session, project and [workload](#).

Justification

1. User consent only matters where there is an ongoing relationship between two entities.
 1. Ongoing could be only for the duration of a session, but current use cases are longer than that.
 2. The penultimate relationship in the digital age is between a user's agent and a resource server or perhaps an AI Agent.
2. Before the user (via their agent) consents to sharing, the requestor MUST display credentials of its own so the user can understand the identity of it (in the full meaning of identity)

1. This is a brand relationship (as opposed to a legal entity identity) for all of the users that are targeted by IDEF profiles.
3. **Consent to Create Binding** (registration of a user) is required before meaningful consent can be given, because it is only after registration that the user can expect to receive notification from the service.
 1. Governmental legal notices could side step this requirement, but that is really ineffectual from a user's perspective.
 2. All such notifications should be filtered by the user's personal AI.
4. The first and mandatory consent granted by the user during registration is some sort of (1) identifier (2) authentication method and (3) email or sms phone number. (these could all be pseudonymous).
 1. If the connection is digital there must be enough of an identifier for the other party to call back to the originator. On the internet it is the IP address. If you are at home without a VPN that is pretty specific.
 2. Often other claims are requested by the resource service during registration, these other claims should be in accord with legal regulations.
5. Once a relationship is established other claims can be requested by the resource server, and optionally granted by the user via the agent. It is this stage that most other specifications address.
 1. Permissions MUST be user directed rather than legally directed, which seems to be consonant with the spirit of the GDPR and other privacy regulations that discuss usability.

Context

In an environment where a subject is requesting the establishment of a binding between its private key and a Provider of any identifier services, the implicit assumption has been that the action of the subject on the server is sufficient. In today's world of gathering a subject's most private information some better means of capturing subject consent is urgently needed. Note that one person may have multiple personas that can be selected depending on the context of the relationship.

Note that the first use case for this message is in the [Health Care Profile](#).

Existing Methods

1. While it is true that methods exist for individual subjects to acquire a certificate for signing emails and receiving encrypted email, the adoption of that method outside of the enterprise has essentially failed and will not be considered as a paradigm for this effort.

2. The most common request today is for an SSL or EV certificate from a Certificate Authority (CA) which works reasonably well for what it is intended to do. While it is possible to set up a CA of your own, we will address the more common case of a CA that has been approved by the major browser vendors. Before the process begins the user selects a Distinguished Name for the site based on the rules established by the CA/B forum.
3. A CSR is an encoded file that provides you with a standardized way to send the Certificate Authority your public key as well as some information that identifies your company and domain name. When you generate a CSR, most server software asks for the following information: common name (e.g., www.example.com), organization name and location (country, state/province, city/town), key type (typically RSA), and key size (2048-bit minimum).
4. An Alternation proposal is to establish a context that can be reloaded to re-establish a prior project's status. For example a [ZeroDB proposal from Toby Morning](#) that works with MCP.

NIST levels of Assurance

- [The NIST 800-63-3 document site](#)
- [Levels of assurance from NIST 800-63-3 are described here](#)

External Definitions used here

- Credential Service Provider (NIST) = A trusted entity that issues or **registers subscriber authenticators** and **issues electronic credentials to subscribers**. A CSP may be an independent third party or issue credentials for its own use. While the CSP is expected to assure the identity of the subject, this page assumes that the identity proofing occurs in a separate service. In these respects the CSP is simply viewed as a function with the physical realization and co-location of services being entirely up to the implementation.

Problems

- Exploits against highly sensitive or valuable content on the web. (aka ex-filtration of valuable content hosted on a resource server.)

Use Cases

- [Remote Attestation Use Case](#) shows the steps that a user can follow to use the [Phone as Health Care Credential](#) with a NIST AAL2 (authentication assurance level 2).

Lifecycle Termination

The question here is how long the user information and authority will be retained in the memory of the recipient. For example, biometric data might be acquired from the user for identity proofing, but not retained for use on subsequent transactions.

1. User Data acquired from the user or other sources, whether sanctioned by the user or not.
2. User Data retained after the transaction is complete.

There are several time scopes that might be of interest. Note that based on the purpose (aka context) different user attributes may be retained for different time frames.

1. Session based - data is not retained after the HTTPS session is terminated.
2. Cookie based - data is not retained by the recipient, but it is included on the user's device and will be returned to the recipient if the user reconnects.
3. Transaction based - data is retained until the transaction is complete. This can be short or long-term depending on the life of the transactions. For example, a 2-year warrantee might enable some data to be retained for 2 years.
4. Relationship based - data is retained until the user terminates the relationship.
5. Legal based - data is retained as long as the law determines.

Termination of a relationship may expire after a certain time, or be revoked by the subject (or agent) at any time. The impact of this on data is complicated as noted above.

Solution

This particular use case relates to the interchange between the subject and a trust authority. A similar interchange would be used between the subject and a relying party.

Identification

It is most desirable that no personally identifiable information is passed from the user to the Agent/Service provider. What is required is a relationship ID, which can just be a UUID in most cases. This UUID will last for the duration of the relationship. When the relationship terminates, the UUID will no longer have any value. It is assumed here that the subject will create the UUID and the agent/service will confirm or reject it. It is possible that the user will select an ID created by a federation service Issuer in some cases in spite of the obvious privacy concerns created by any reusable ID.

An alternate form of identifier would be a DID which is attached to some resolver. The downside to the DID is the immutable ledger which may be acceptable for some uses. The advantage to the DID is the immutable ledger which provides a deterministic audit trail. It is assumed the use of the DID would be frozen when the relationship ended.

Registration Ceremony

- A part of establishing a relationship over some time period between two entities requires that each be known to the other in some meaningful sense. In the case of the service provider, there are multiple solutions which started with HTTPS and are currently undergoing a process of improvement. In the case of the user the needs for identification and authentication proofing vary substantially based on the context. One of the level two assurances is defined in the companion [Distribute Assurance Specification](#).
- At the end of the registration ceremony, it is expected that a notification channel will be opened from the service provider to the user to meet the requirements of the GDPR or CCPR.
- In general registration will require the ability to notify the user of changes, breaches, etc. That capability requires information from the user which is supplied in a **Consent to Create Binding**.
- Typically a service provider will request additional information from the user for its own purposes. That should not be confused with the necessity for a notification binding as is described below.

Formal Methods

Typically AI prompts start with simple queries to some AI that are gradually refined to get to the desired results. For more involved projects it is better to start with some formal statement of the problem to be solved, or the result desired. For example a table of requirements or a knowledge graph. For a long-lived service the most common model is with a data flow diagram.

In generation this can be considered to be a formal definition of a prompt engineering process that is constructed in advance of making the first request to any AI agent.

Consent to Create Binding Request

The following is the current understanding of what needs to be included in a Consent for Binding Request. This set of requirements can be met by the format of the [OpenID Connect 1.0](#) id_token. The security considerations might be somewhat different than that spec. Also the id_token is viewed as the final step of an authentication interchange; while this message includes a request for a response from the target.

1. Identifier of the message

Protocol dependent - may have message identifier, version number, unique sequence number, nonce or state.

2. Issuer (aka Organization or Audience)

MANDATORY - URI OF the CSP (Audience of the request). This would be the entity that supplies a public key when encryption is required. Rather oddly, the issuer is the target of the message. That could be made less confusing.

3. Subject (aka Grantor)

MANDATORY - this is the identifier from the user that will be the subject of the binding. It is called a UUID above for purposes of discussion but could be in any format the user chooses so long as it is unique to the server. It serves nearly the same purpose as the DN of the X.509 certificate. Whether this subject identifier has any binding to a real-world entity (like a human being or other legal entity) is to be determined by the purposes to which the resulting entity statement will be put.

4. Subject Role

MANDATORY - this is the relationship between the subject and the owner of the information (e.g. patient). Typically will be self or guardian. Note that this sort of delegation could be a reason for creating a slightly different version of this document to handle delegation.

5. Context (aka trust framework, may be multiple)

MANDATORY - The trust authority that establishes the context. This authority could be as small as a family or as large as a country.

6. Subject Attributes (aka Claims about Contact and other information)

OPTIONAL - although a contact email or mobile number might be a requirement during registration for notices.

7. Permissions

MANDATORY - applies to above subject information. This is purely opt-in. No reference denotes no consent. Note that the access permissions granted may be part of this message or may be included by reference with an Authorization Code for access. There may be a language like cedar to encode permissions.

Other contractual terms may be added (like MyTerms) that will form a contractual binding if accepted by each service. It is projected that this contractual binding will be legally acceptable in most jurisdictions.

8. Device Statement

MANDATORY for AAL2 and higher certifications. It is used by the Issuer to verify the level of protection provided to the Private key of the certificate.

9. Identity Proof

MANDATORY - although this request might pass [Responsibility](#) for proofing to the CSP if it is willing and able to provide proofing for the subject. The content of the proof might need to meet specific requirements in IAL2 and higher assurance credentials. This is an array so that more than one proofing document may be included. Note that the legal jurisdiction or address of the subject identifier may be part of this structure. Note that a

request for proof that the user is not a robot (CAPTCHA or personhood) would fall into this category. "None" would be an acceptable value for the element.

10. Purpose of Use (aka reason or intent)

MANDATORY for any level of assurance greater than level one in any of the 3 categories of assurance. See the [Purpose Query](#).

11. ACR

OPTIONAL - this is only useful in the case where the Purpose is not adequate to establish the required levels of assurance of the resulting Entity Statement.

12. Issue Date

MANDATORY - Linux epoch date is default

13. Expiration Date

OPTIONAL - Linux epoch date is default

14. Subject Public Key

MANDATORY - must be included directly or by reference.

15. Signature

MANDATORY - using the above key.

16. Encryption

MANDATORY - but it can be provided by HTTPS - Note that the above items are all a part of the JWS - or signed JWT. Encryption exists above this JWS content and may include multiple JWT elements in a JSON object.

Consent Receipt with Binding

The Binding can consist of any sort of authorization or certificate that shows that binding of the user to the resources has been created or reconfirmed.

- n.b. The US Health care proposal has decided to separate the binding and consent receipt into two separate JWT/JWS signed Jose documents that could be combined into a single JWT/JWE Jose envelope for encryption.

1. Binding Statement (aka [Consent Receipt](#) or [Entity Statement](#))

Response to the REQUEST - this is the name, version or identity for a message returned by the issuer after the process has been completed. It is then made available to any legitimate request. It is signed by a well-known key belonging to the issuer.

2. Issuer (of this receipt aka Organization)

MANDATORY This the entity that the subject is forming a binding with. It may be a brand name rather than a legal name.

3. Role of Issuer

MANDATORY - this is the role(s) that the issuer takes. (eg CSP, Identifier Provider, Data Controller)

4. Legal Name

OPTIONAL - this may be included if the Issuer is not a legal entity. It is required that a legal entity be responsible for the receipt.

5. Subject

MANDATORY - this is the identifier supplied by the user. This receipt is bound to this identifier.

6. Claims

OPTIONAL - this will be either the categories of permissions granted, or the permissions themselves.

7. Issue Date

MANDATORY - Linux epoch date is default

8. Expiration Date

OPTIONAL - Linux epoch date is default

9. Authority Hints

OPTIONAL - List of federation authorities subscribed to by this issuer

10. Issuer's Public Key

MANDATORY - must be included directly or by reference.

11. Signature

MANDATORY - using the above key.

12. Encryption

MANDATORY??? - using the Subject's key.

Comparison with a representational Mapping

1. Kurt's diagnosis is right, but incomplete

Kurt Cagle described how we are experiencing a representational crisis, where the maps (institutions, metrics, categories, policies) no longer correspond to a fast-moving, AI-driven territory. This mismatch explains why interventions arrive too late, too small, or are captured before they work. Where his framework stops short is here:

- It assumes the solution is a **better map**
- A more dynamic ontology
- Better modeling of consequences

But in an AI-agent world, the core problem is not that representations lag reality. It's that representation itself is no longer the primary organizing layer.

2. The shift intelligent agents force: from representation → relationship

Traditional systems govern by **description**:

- roles
- classifications
- metrics
- thresholds
- policies

These work when systems are:

- slow
- decomposable
- centrally legible

AI agents break all three. Agents act continuously, adaptively, and across contexts. You cannot pre-represent their future states in the way Kurt's "better map" approach still implicitly assumes. This is exactly the problem the "map vs territory" metaphor runs into when the territory is now co-produced by the map itself, a concern echoed in broader AI literature on abstraction failure.

A **relationship-based paradigm** starts from a different primitive:

Not *What is the system?*

but *Who is in relationship with whom, under what obligations, with what recourse?*

3. Why relationships outperform representations for agents

A. Relationships bind *behavior*, not just description

Maps describe, Relationships **constrain**.

A relation encodes:

- mutual commitments
- expectations
- escalation paths
- consequences attached to breach

In human systems, we already rely on this:

- fiduciary duty
- professional responsibility
- legal personhood
- contractual obligation

Agents do not need full autonomy to be governed relationally. They need:

- **assigned roles**
- **explicit counterparties**
- **traceable accountability chains**

This addresses Kurt's "consequence differential" directly: relationships collapse the distance between action and responsibility, instead of trying to represent consequences abstractly.

B. Relationships are resilient under change; representations are brittle

Representations fail when:

- categories drift
- incentives realign
- edge cases dominate (which is *normal* in AI systems)

Relationships persist even when:

- tasks change
- tools evolve
- contexts shift

This is why governance systems built on roles (doctor–patient, trustee–beneficiary) survive technological upheaval better than those built on fixed categories. No new map is required every time the environment shifts; the relationship carries forward.

C. Relationship paradigms scale with agency

As agents become:

- more autonomous
- more composable
- more capable of delegation

Representational governance explodes in complexity. The number of states to model goes superlinear.

Relationships scale *linearly*. (Trafton 2026)

- each agent has defined obligations
- every action has an accountable anchor
- escalation is relational, not systemic
- Human relationships are seldom symmetric
- There is seldom one way to know in advance who is in control
- Agent relationship will likewise be tilted by control

This is a fundamental asymmetry Kurt's framework does not fully exploit.

4. What changes concretely under a relationship paradigm

Under a representational approach:

- we ask whether the **model complies**
- whether the **output matches policy**
- whether the **classification is correct**

Under a relationship paradigm:

- we ask **who vouches for this agent**
- **who bears downstream responsibility**
- **what happens when trust is violated**
- **how redress is triggered**

These are not semantic questions. They are operational ones.

Importantly: **you do not need perfect foresight** to govern via relationships. You only need:

- clarity of obligation
- enforceability
- continuity of accountability

This is why relational governance has historically handled uncertainty better than technocratic planning—a lesson implicit in many failed attempts to “map” complex political-economic systems.

5. Why this matters *now*, not later

Kurt is correct that we're governing a 21st-century economy with 20th-century tools. Where the relationship paradigm advances the argument is in showing:

- You **cannot** out-map a system of adaptive agents
- You **can** anchor such a system in durable relationships
- Governance failures today are less about bad intent than about **misplaced primitives**

The crisis is not just representational.

It is **ontological**: we're using the wrong unit of organization.

6. Summary

Kurt asks how to build a better map for a moving territory;
a relationship paradigm accepts that the territory moves and governs by binding actors
to one another rather than to static representations.

References

[Anne Trafton 2026-06-11 "Would you return a favor? Scientists say it depends on the relationship" MIT News](#)