UMA legal subgroup notes

Standing agenda for 2019: Work on producing our second legal-business framework report by September, initially focusing the work on use cases that illustrate each of our mappings from business relationships (and changes in those relationships) to UMA technical artifacts.

UMA legal subgroup notes

Date, Time, Documents

"Legal" topics are currently being covered in a separate legal-business framework meeting series. See the UMA calendar for details. Current documents being worked on (WG participants will receive edit access for the asking; others may receive view access):

- Business-Legal Framework and Use Cases (GDoc)
- Mapping Graphics (GSlides)
- Legal Role Definitions (GSlides)
- Definitions and Use Cases Spreadsheet (GSheet)

2020-01-28

Attending: Eve, Domenico, Colin, Tim, Lisa

This is the last time we'll be meeting at this hour where it's the "biz-legal call". We'll see what the results of the Doodle poll are in terms of a new hour. Please fill out that poll!

Tim likes that the Guardianship paper really wrestles with the temporal dynamics of guardianship. Tim had introduced the concept of "diachronic" issues in our work on the first Report, though it looks like the word itself didn't survive the editing process through to the final version. These aspects reflect (literal) life cycle changes that traditional IAM typically handles through workflow approvals and the like. Colin would have liked to see a reflection of existing work. Likely people are not fully understanding both UMA and OAuth, both of which handle delegation of authorization in some subtle ways. We need to make this really easy to understand with some demos.

The very first step in a Me2B relationship, delegating a guardian – or other RRA (resource rights administrator) – relationship, is not yet today interoperable in an OAuth or UMA world. Does profiling OAuth with a claim representing this semantic make sense?

We looked at the original NZ POC case study to see if it had any examples of guardianship or interesting delegation. The Aroha/Bailey use case is about "classic" delegation of access, though it includes mobile notifications of IoT data, which is nice. The ministry of education/picking up kids use case is about chaining delegation.

At Kantara we produce reports and specifications. We could analyze the use cases covered by the paper; our BLT work address guardianship along with various non-guardianship delegation use cases between data subjects and RRAs. UMA's architecture and Sovrin's architecture are pretty different. However, Identos has created an extension that seems to be potentially valuable in privacy-enhancing an UMA AS in an SSI-ish context, and Adrian and his cohorts have integrated uPort for providing RqP verifiable claims. We'd like to get a broader shared understanding in the group about these options.

2020-01-21

Attending: Eve, Andi, Domenico, Lisa, Nancy, Tim, Colin

Through her colleagues, Eve recently came upon the work of the Sovrin Guardianship group (which overlaps in some respects with our biz-legal work) and the DIDauthz work (which references OAuth and UMA). Nancy brings up the FAST work in healthcare, which is trying to accelerate solutions to common challenges. What is the best way for us to accelerate our own work and goals, and how we can center on the end-user perspective (the "User" in UMA)? Lisa notes the DIDUX working group, given this desired perspective.

Tim notes the rationale stated in the paper on p. 9 for the guardianship work: "In short, carefully constructed guardianship is essential to SSI. Without it, SSI solutions will either tend towards centralisation or exclude billions of people." This is somewhat a weird way to think, in Eve's opinion, because it's constructed around the goal of digital identity (an "input metric") as opposed to what people actually want out of digital services (an "output metric"). If we don't get out of this mindset, we may not get to the point of being creative enough to solve the next generation of problems. Tim notes that the UN and the World Bank are saying that legal identity is a human right. (Though "legal identity" is not the same thing as "(a) digital identity".) She takes the sentiment of wanting to solve the thorny problem of "offline and online" mixes of people, though. Lisa suggests that we pull Adrian into an analysis. Eve thinks Justin could shed light too.

Eve is going to pare down our meetings so that we only meet on Tuesdays. She'll send a note to the list just to be sure this doesn't cause any heartburn for our voting participants.

Homework for next time: Everyone please study the paper and the use cases, paying particular attention to the terminology and concepts, and seeing if we have comments or could use some of them. Eve will look into her team's ability to develop demos of the use cases in the paper (or equivalent), on an UMA /OAuth basis and potentially SSI. Others are welcome to do that too.

We are meeting next Tuesday. No meeting on Feb 4 as Eve is traveling.

2020-01-07

Attending: Eve, Tim, Lisa, Domenico

Could UMA be additive as a solution to the human rights efforts such as in the UN? Digital human rights is a driving motivation for Lisa. Rights and corresponding responsibilities need to be ensconced in technology where possible. An old writeup about UMA implications of Privacy by Design principles (for which the tiny URL is http://tinyurl.com/umapbd) is probably a bit out of date but helpful on this score. Me2B recently did a small ethnographic study, discovering some interesting – if perhaps nominally obvious – things about surveillance and awareness about it. There is a faction among healthcare IT people saying that informed consent is literally impossible. This is akin to the *qualia* problem: you can't inspect someone else's brain to know what they know. If this is the case, and even given that "consent is stupid", it's better to *assert permissions* as proposed in the LeVasseur/Maler paper. If people won't even read books now because their attention spans are shot in the modern era, then how can we expect them to read ToS? There's a Japanese word for the pile of books you have but haven't read, tsundoku, and now we have the digital version too.

To bring the news about what UMA can solve and how it can support "IRM" challenges, we want to publish the material we've put together in smaller bites, possibly in blog posts.

We talked a bit about the JLINC protocol, which can be found here.

2019-12-17

Attending: Eve, Cigdem, Domenico, Lisa, Colin

Note that we did rename (renumber) the states, so that the "little Johnny" use case is now State 2 because it's the next most common.

Domenico shared a great new non-tabular graphic that is helping us think about what we really mean by the first four states. (See the wiki version of these notes for the graphic.)
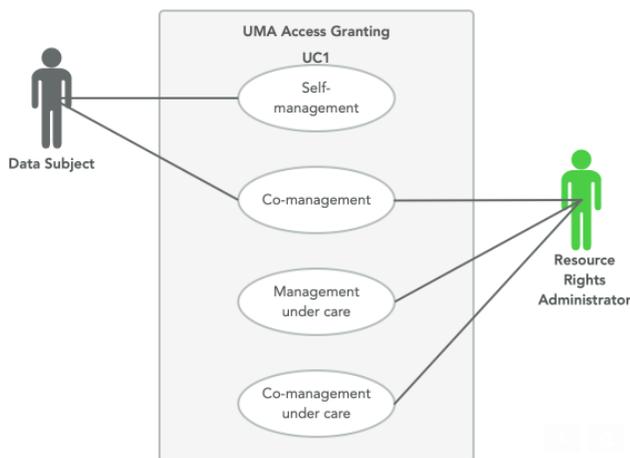
In the "Mapping graphics" deck, Eve had put some comments about needing to see the DS-as-RRA aspect as donning a role, and Cigdem raised this as well. The states we've been talking about are at a pretty operational level, meaning they allow us to specify details of implementation such as role changes and provisioning/deprovisioning requirements that could built into compound workflows. We're likely to have to go to the level of tackling the "multiple Data Subjects" use cases.

In healthcare, you could commonly have state transitions like 2-4-1, or 2-4-2-1 as little Johnny has his mother added to manage his EHR, then his father, then grows up and starts managing his own resources. We actually map cradle-to-grave scenarios somewhat like this in the "Legal role definitions" deck.

How does one map three-dimensional values in two-dimensional space? You split out one of the dimensions to two flat ones, or you have multiple tables. Ugh.

We observe that adding and removing RRAs is a provisioning/deprovisioning action, and adjusting whether the DS is one of the RRAs is a role change action, what if we think of this as two different "tables" (maybe tables in a technical viewpoint, but potentially not for the document)? They're orthogonal actions, which could be built up to form a "workflow" that triggers when you need to move from one state to another, like when Johnny e.g. hits a certain age, or succeeds in asking the court to be emancipated from his parents, or when a couple divorces, or when a temperature sensor hits a certain temperature, etc.

There can be harder edge cases for what to do around, say, historical bank account data vs. new data going forward in the case of multiple-down-to-single bank account data. What do banks do today? Do they close the joint account and open a new single-owner account for the newly divorced person?

Lisa shared a great "IRL mapping sketch" that is inspiring us to add a graphical version of the "typical use case" to the doc.

2019-12-10

Attending: Eve, Andi, Cigdem, Tim, Colin, Mark, Nancy

Lisa and Eve subsequently discussed the intro section and Eve did some more editing of it. This led her to bring up a few more legal party terminology and definition questions. It looks like we can remove the "Individual or Legal Person" phrases where they occur currently in the definitions of RRA and RqP. Let's do that. (Eve edited this live on the call, in both the canonical spreadsheet version and in the report.)

We discussed the question of Requesting Agent/Requesting Party. Tim made the excellent point that, until we hear that the outside world has a problem with the terms, we probably don't want to obsess any more about it. We could change it later if it turns out to present friction. We acknowledge that "agent" is a very different thing in the legal and technical worlds. Capitalized words are used in their legal party senses and we say that in the report, so legal experts and similar should be prepared to understand such terms in their legal senses.

We discussed whether to cram mentions of "agency contract" and "access contract" into the pentagram diagram (new nickname!). Should we do "progressive disclosure" in the document and have a version that has just the dashed pentagram, possibly with the agency and access contract wording added, and then a version with the delegation and license details added? If he is willing, let's ask Domenico to create a short series of diagrams.

AI: Eve/Domenico: Eve to ask Domenico to create several pentagrams (these are all additive):

- One with just the dashed pentagram lines and agency contract/access contract labels (as illustrated on the old "spaghetti" diagram on slide 7 here) – delegation/licensing relationship arrows removed
- One with all the delegation/licensing relationship arrows added back, as currently exist in the diagram
- One with a "spur on the left side with Data Subject-to-RRA relationship arrows added (as illustrated on the left side of slide 9 here)
- One with a "spur" on the right side with a Requesting Party-to-Requesting Agent arrow added (as illustrated on the right side of slide 9 here)

AI: Eve/Domenico: Eve to ask Domenico to create two table-oriented "state" diagrams (slides 3 and 4 here):

- One without arrows
- One to reflect the state changes with the arrows

2019-12-03

Attending: Eve, Domenico, Lisa, Nancy (regrets: Tim)

Lisa's and Eve's "Beyond Consent" article is now in a near-final version; Lisa kindly agreed (when asked by Eve) to distribute this version to the UMA WG.

Consensus: We looked at the Typical Use Case section of the doc and concluded that we do like having the "typical use case" for technical roles presented, followed by the formal definitions of said roles. Let's also plan to publish the spreadsheet as a companion Report because it can serve as the "source of truth" of definitions and relationships. We can prepare an Annex or Appendix or whatever we want to call it in this Report that consists of tables, along with a citation to the spreadsheet that has "live" data. We can say "In case of any discrepancy, the spreadsheet has the most accurate data" etc.

Assumption: We note here – but don't want to note in the Report, because it would complicate things for the readers – that EHR data typically gets selectively copied to a portal, and doesn't "live" there. If we elucidated all of our assumptions and caveats in the "typical use case", we'd never get anywhere. Nancy is clarifying for our group exactly how SMEs would split hairs on primary storage of health data. We agree that our audience will be motivated simply to understand the UMA paradigm at a relatively high level at this stage.

Lisa asks (paraphrased): When we say the resource owner "manages" resources at the RS and "controls" access at the AS, is she doing that at an app of some sort? Why are we saying Alice is interacting directly with services here and not apps? She does use apps, but we don't talk about them in protocol-related conversations because those interactions aren't "on the wire"; there are no UMA-dictated messages involved. However, those who deploy UMA-enabled systems do have to concern themselves with the implementation and deployment of such apps! Eve's recent webinar with Steve Giovannetti covered the four elements of integration and deployment needed:

- Deploy an AS. We have noted many times that policy engine functionality is something each AS can compete on because it lets ROs get more sophistication in policy-setting ability. However, often in healthcare use case conversations, the idea of policy interoperability comes up too.
- Integrate/create your client app (if you already have a client, you have to UMA-enable it) – is there room for creating "UMAlet" middleware? Gluu has talked about its open-source middleware in the past.
- Deploy your RS (if you already have RS apps, you need to UMA-enable them) – there are a couple of different approaches, gateway/proxy (could be in the cloud) and an SDK for UMA-enabling apps. The former seems preferred. Eve calls these "protectlets" 🙂.
- Encourage data partners to implement their own RS's – if you have a wide data ecosystem (i.e., multiple RS's in different domains), your partners will have to UMA-enable their hosts. Helping them do this is good business! See just above. Nancy notes that Patient Centric Solutions' new PatientShare offering does exactly this.

This whole list opens up the idea of publishing an auxiliary document called something like an **UMA Deployment Guide** and stimulating open source (once more) to encourage deployments.

Proposal: Lisa suggests that, since the Operators are extremely likely to be Legal Persons, we should acknowledge the obvious and enhance the formal definitions in the document somehow so that it won't trip up people who read them there. Maybe we can insert some bracketed text and/or italicize some? She will draft some candidate text and let's discuss it on the list. There are so many considerations around this question – rhetorical, UMA maturity, etc., that it's not even funny.

Domenico comments on the concept of a "data ecosystem". We've been talking about wide ecosystems for a while, but it seems there are different kinds of wide ecosystems. Lisa says: "UMA is an ecosystem builder." Nice! Eve has often called OAuth2 a protocol framework because it has so many choices that it's not just one protocol; unfortunately that has risked interoperability and security. XYZ is trying to enable protocol-building without all those sacrifices. UMA is meant to make it possible to build data – and data-sharing – ecosystems without sacrifices and compromises.

2019-11-05

Attending: Eve, Nancy, Tim (regrets: Cigdem, Andi)

We have been dealing with the "multiple RRAs" state table as if there is always and only a single DS. For simplicity, perhaps we should continue in this vein. It would be insanely complex to imagine that there are multiple DS's as well. The "co-equal rights" question is about resource administration rights specifically. However, a number of real-life use cases have the "multiple DS" situation: joint bank accounts where both account holders should have exactly equal rights, etc. In implementation, this is a problem and is generally handled badly, in that one person is made "primary" and the other(s) "secondary. In some use cases this divide is appropriate (main householder vs. other people in the household for digital streaming services), but for things like a joint bank account or other spouse-type membership accounts, both/all holders should really be co-equal – until such time as there is a divorce or other split and the number of DS's goes back down to exactly 1. What if two brothers share a bank account and one of them has to be removed from managing it for malfeasance and the other brother is managing it for the two of them, somewhat like in the Gravity Payments case?

To clarify whether we need to focus on multiple-DS use cases and whether they are sufficiently valuable to solve, Eve will ask an FSI/bank expert to join us at our next meeting. Nancy will do the same for a genomic data expert.

The current use cases that we've mentioned under State 3 in the table ("joint bank account each with equal access, joint tenancy") look like they're not actually correct because they probably refer to multiple DS's, so we have to sort out whether we're going to treat multiple DS's or not, and move the use cases over if so.

AI: Nancy and Eve: Find experts to join our next call to establish existence and priority of multiple-DS use cases in different sectors.

2019-10-29

Attending: Eve, Cigdem, Nancy, Andi, Adrian, Tim (regrets: Domenico)

We worked through the section after the Introduction. We want to get to the RRA-centered states and state changes content next time. Nancy has provided some rough content/use cases that are healthcare-focused that we can use as a base. Let's also add (say) FSI-focused use cases. There is also the co-equal rights challenge (multiple co-equal resource owners/RRAs? and also DS's?).

2019-10-22

Attending: Eve, Colin, Nancy, Andi, Domenico

We worked directly in the document, accepting suggestions and editing the results.

We discussed whether we should keep "Alice" for consistency as the resource owner name, vs. use variety (e.g. the "Jane" example as in Andi's new "simple example up top).

AI: Domenico: In the Mapping Graphics slide 19 graphic, can he please ensure that the "technical terms" (currently with no-color background outside the legal terms) be distinguished somehow (maybe with italic or with a color background) and be all lowercase?

In the Typical Use Case section, let's try telling a brief narrative story about the life insurance use case, either before or after the technical definitions, so as to make it easy for nontechnical readers to understand those definitions. We do need the technical definitions – otherwise we can't delve into the permission token/license work that the paper needs to do.

2019-10-01

Attending: Eve, Andi, Domenico, Nancy, Tim, Cigdem

If we are addressing a business-legal audience, should we start with a relatively technical approach, or should we start with use cases? We do have a challenge with people not understanding UMA. We are finding confusion out there. We did say at the end of the last paper: "The next papers in this series will explore **the application of UMA licensing model to specific use cases** for various categories of owners and custodians of personal digital assets in the global digital information society." So we could start with a (full?) set of concrete use cases, and then work through them step-wise.

So let's have an outline like this:

- Part 1: Introduction
    - Explain UMA at a high level to a primarily business-technical audience
        - Explain technical terms in the context of State 1: RO = DS = singular RRA
- Part 2: Delegation of Resource Rights Administration Use Cases
    - Use cases written in simple language (maybe one paragraph each), illustrating various of the DS/RO/RRA states and state changes
    - A state machine showing the states and changes
    - A section working through those use cases with the permission token flows and consequences
- Part 3: Sharing relationship use cases (policies change based on changes in relationship with RqA?)
    - Subsections are like part 2's subsections...
- Part 4: Delegation of Requesting Agent Role Use Cases
    - (e.g., Alice shares her connected car with Bob and he gives his keys to her car to the valet for parking)

It's perfectly fine to blow up the text that's in the current document vs. trying to make the existing text work. Nothing in there is sacred. Andi's schedule has now cleared enough that he should be able to contribute significantly for the next couple of weeks.

Future meetings: The call next week (**Oct 8**) is already cancelled. Eve can't make the call on **Oct 15** but Ruth P has kindly agreed to start the call!

2019-09-24

Attending: Eve, Cigdem, Domenico, Nancy, Tim, Vlad, Sal

At the DS-RRA wrangling layer where UMA doesn't have any artifacts, there will typically be a lot of identity and access management taking place. What technical artifacts can we expect to be used there? Typically many proprietary ones, but also potentially some standardized ones, such as federation standards like SAML and OIDC, maybe provisioning standards like SCIM, maybe auditing-friendly standards like consent receipts, etc. We have already mentioned consent receipts in our "devices and artifacts" capture.

The topic discussed last time was: Could/How could UMA be used to achieve user control of the initial user relationship with a service provider? The answer in today's world is that UMA was designed with today's imperfect situation in mind, which presumes opt-in cookie consent, opt-in terms and conditions, opt-in OAuth app connections, and opt-in AS-RS connections for UMA, which are all sub-optimal. However, the Lisa/Eve paper proposes an architectural way forward for fixing these broken patterns, which is to use the "Open Banking trick" for intent registration. The trick uses the nexus of the Open Banking APIs and OIDC request objects to allow transaction authorization of payment of a specific amount of money for a single payment, vs. authorization of an OAuth scope for an indefinite period. Theoretically, this "intent registration" ahead of time by a user at a client app (TPP) could be used to allow transactional (one-time) "intent registration" of user-centric terms and conditions, not just payment of money. Flows still to be worked out for each use case of user/service interaction, of course!

Cigdem had commented about moving up the Legal Parties discussion in the document and we agree. She will do some "invasive" editing of the doc.

Eve can't make Oct 8th or Oct 15th. It sounds like we have critical mass of people to join on the 15th so Eve will ensure someone can open the bridge on that day. In the meantime Cigdem will edit the doc.

We haven't discussed Nancy's latest submitted set of use cases, including healthcare and banking. One of Eve's goals is to include as many submitted real-life use cases as possible. We've already got quite a few. She's started a couple of templates for some more in the "mapping" slide deck.

2019-09-17

Attending: Tim R, Lisa L, David Thelander, Pete Palmer, Ruth P (Kantara)

The group discussed whether and how UMA currently treats first order legal relationships (i.e. the creation of the initial legal relationship between the consumer/data subject/legal representative and the service provider and/or resource server operator. For purposes of this first order legal relationship, the group further discussed the possibility of deploying a 'reverse EULA' as discussed in the recent paper published by Lisa and Eve. Lisa reported that the Me2B organization has a legal subcommittee that is now actively developing model 'reverse EULA' language. It was suggested that UMA might collaborate with Me2B on this effort. [Editor's note: The Proposed UMA Business Licensing Model contemplates that the ASO, using an Access Contract vehicle that authorizes licensing, is the entity that creates the first order legal relationship with the Service Providers and Resource Service Operators.]

Follow-up actions: Lisa has offered to share information on the Me2B efforts to create model 'reverse EULA' terms. In particular, Scott David has prepared excellent initial legal materials. David T has legal expertise with software licensing and also offered to assist with developing standard licensing terms.

2019-09-10

Attending: Eve, Andi, Nancy, Domenico, Vlad, Cigdem

Our agenda for today: "Start to document the full value and structure of a final "WG draft report" (or is it a spec?) deliverable that we can put together with all this fodder based on our work over the last several months. We had said our deliverable timing would be September... :)"

The value-add of "UMA-biz-legal-relationships" above and beyond "UMA-technical", we think, is that (e.g.) the operator of a service or client, or the RRA or requesting party, can auditably prove what another party needs it to prove regarding a commitment or promise it has made to the other party. And if necessary, if enforcement can't take place at a technical level, it can take place in a court of law as required (because you have audit logs and non-repudiable artifacts to point to). All of these loosely coupled parties need to cooperate and have their interests sufficiently aligned to have selective sharing work. See examples listed in the Sources of liability tension section on the UMA Legal page here.

Which party cares most about changes, additions, and subtractions in the RRA role? The ASO. The PAT is the "thin reed" of technology (a technical artifact) that binds the ro, as, and rs. Is it enough to give us a technical means of controlling lots of biz-legal-relationships scenarios?

We have several sets of scenarios. One is the ro/RRA state diagram. A second could be a rqp/Requesting Agent/Requesting Party state diagram (that doesn't exist yet). A third that doesn't have a diagram yet is where only policies (a non-UMA artifact) change in response to the ro/RRA sharing on the basis of a relationship. Example: Alice wants to share a set of resources with "whoever is in the 'spouse' relationship" with her. If she indicates that no one is now in that relationship, her policies reconfigure and sharing ends. What's attractive is that both Alice and the ASO and RSO(s), by extension, can all know that sharing ended even if it's a big complicated set of resources that was shared. In this third case, is it really in our scope to be talking about these policies? We could say they're out of scope, but some "consent" (could be described as authorization policy and sometimes is) data formats are standardized, e.g. in HL7 (a Consent record) and XACML, and it is sometimes desired to "federate" consent servers, and even though UMA doesn't standardize policy data formats, it's sort of an adjacent use case. It exists on a continuum with UMA use cases. Plus, all three sets of scenarios seem to end up needing to be combinatorial at the end of the day.

Note that HL7 federated consent implies that you have a federated consent server that you can read and interpret. Consent itself is PHI (personal health information). While UMA technically doesn't require policy conditions to be stored at an AS (this is deployment-specific), it does assume policy-setting can happen at the AS; there are numerous examples in the specs throughout. And the privacy considerations in UMA FedAuthz talk about PII involved in resource registration and policy-setting only implicitly ("the authorization server may come to learn a great deal of detail about Alice's health information just so that she can control access by others to that information"). For example, if Alice tells her AS to share all her data except HIV-related data, that exposes that she may have HIV-related data. The HEART profile for OAuth/FHIR tries to protect against assumptions about existence of confidentiality /sensitivity scopes, but it's still hard not to make assumptions.

Just because you *can* separate the AS and RS, it doesn't mean they're necessarily in separate organizations/domains. Some privacy considerations may not apply in cases where they are run by the same organization.

Let's start drafting collaboratively. A GDoc is a great way to start, following the (Kantara-inflected) RFC style. Andi and Cigdem are willing to jump in on co-authoring with Eve. Nancy is willing to be backup 🙂 after the next two weeks.

AI: Eve: Start a skeleton draft doc for next time.

AI: All: Sort through the next two batches of scenarios and work on the wording we're using in the second scenario batch in our next meeting.

2019-09-03

Attending: Eve, Andi, Domenico, Lisa, Nancy, Peter, Tim, Colin

Eve has put a new table in the Business Model Mapping Graphics deck for our perusal. We discussed ways to improve it, and moved around the arrows to make it more understandable and added use cases.

We posed the question from a few weeks ago "For the care scenarios, do we need a primary carer/super admin type of role or are all entities in RRA equal?" It's not just for "care" scenarios. When a RRA/DS has the individual capacity to grant access to another (say) individual, then they are *not* equal, and access that can be granted can be taken away. If the two are truly equal in their rights (such as joint tenancy or a joint bank account), it's because there are institutional factors at play such as a law granting that equality. If it's because they are married and there's a law saying spouses have totally equal rights in a bank account, and they then get divorced, it's by law that the account access gets disentangled, not just by "access management technology". So if there were two RRAs (two married bank account holders, both DS's), and we went from state 2 to state 1 when then got divorced, it was because one of them is no longer going to be a DS in future and should no longer be an RRA in future. Presumably in this case it's not automatic but manual, and the parties need to tell the system what action to take (e.g. which party wants to remain in the system, when to take action, etc.).

How complex should we get? In this new table, do we assume there's one DS? Recall that last time we discussed multi-DS scenarios and it rapidly became so complex we became uncertain about how to capture it.

What are we documenting? The value is to be able to show proof as an ASO that the correct RRA is doing the sharing. Each "design pattern" represents perhaps many use cases in many different sectors.

We posed the question from a few weeks ago "Federated authorization or not? Arguably, personas of a DS and a single DS have the same identity verification risk. If personas, then "DS" in the state machine is replaced with "persona". The key assumption: KYC works – if that fails, all the rest fails." See the new diagram in the Biz-Legal Mapping slides; any one resource owner/RRA is going to have a single login at an AS but could have a unique login at each RS that is federated with that AS. The single AS login could be strongly proofed, meaning that the ASO could prove strongly who that RRA is.

Eve has moved our biz-legal meeting **a half-hour earlier** next week as she has a partial conflict.

2019-08-06

Attending: Eve, Adrian, Cigdem, Domenico, Lisa, Nancy, Sal, Thomas, Vlad, Tim, Colin

Cigdem sent some thinking around the "multiple DS" cases:

Regarding multiple DS'es the table would look like this.
The Multiple DS case may not occur often, but is present in the case of a shared bank account between partners.
When a couple opens a bank account, they may start at State 5, when one of the couple goes under care, they transition to State 6,
When one of them dies, the end state is State 1 if the other partner is the sole beneficiary, or if there is another person, it may go back to State 5, or State 7 etc.

Not sure how to represent this technically – it seems to work for DS=2; not sure it extends to a larger number.
Do we have cases where DS is greater than 2.

```
DS>1 | Rep>0 | Carer>0 | State
 F   |  F    |  F      | State 1: Self-management (Single DS)
 F   |  F    |  T      | State 2: Management under care (Single DS)
 F   |  T    |  F      | State 3: Co-management (Single DS)
 F   |  T    |  T      | State 4: Co-management under care (Single DS)
 T   |  F    |  F      | State 5: Self-management (Multiple DSes)
 T   |  F    |  T      | State 6: Some/All DSes under care (Multiple DSes)
 T   |  T    |  F      | State 7: Some/All DSes co-manage with Reps (Multiple DSes)
 T   |  T    |  T      | State 8: Some or all DSes under care and co-managed (Multiple DSes)
```

If we're in State 5, but I decide to co-manage my bank account with another party, what does my partner get to say about the situation given that they are a DS too? What if the two partners have different accountants; does the second partner get to/have to consent to the first partner's sharing with a chosen accountant? Divorced parents may have different custody rights and agendas – but this may fall under the single-DS case. Joint bank accounts and DNA data seem to fall under the "multiple-DS" case. What use cases are truly realistic around the RRA and sharing breakdowns beyond that? If there are whole state machine branches that are truly not realistic, let's not map them. Are trusts another example of potentially multi-DS and multi-RRA? It seems so. Trusts generally are for administering physical assets, but if they could be used for virtual assets as well, then what we're doing here could provide real added value.

The ASO/AS role provides value in terms of the actual permissions as they change. Right now we're focusing on the RRA/RO role changes, which are at the "head" of UMA-technical. Is there another service/party needed, something like a "relationship manager", to manage the state transition stuff? In identity management, there are some solutions that do things like this under the heading of identity lifecycle management, relationship management, and essentially "multi-identity lifecycle management".

Thinking out loud: Three siblings open a shared account, A, B, and C. They appoint an account manager, D, as an RRA. So each is in state 6. Brother B decides he wants to self-manage the account, so he may need to get permission from the others before transitioning to state 1.

Sal and Mark at OpenConsent have been doing some work on policy states and state transitions; we will hope for a demonstration soon. We think it's complementary.

Some concrete scenarios: This PC Magazine article explains a) how to delete your Facebook account, and also b) how to request removal of accounts for someone who is medically incapacitated and c) how to get an account memorialized when someone dies. The latter two seem RUFADAA-like and have very close analogues in our Cradle to Grave use cases. For example, in use case (b), the party making the request has to prove they have the right to be the RRA (a Rep). In use case (c), the RRA role is self-asserted but they have to prove the DS has died.

Should we think of multiple state machines, one for every DS, and not maintain States 5-8 as formal constructs? Does the "relationship manager" center on resources first (e.g. the bank account) and then look up the relevant current DS(es) and RRA(s)? Relationships tend to be graph-like, without a single "head of the hierarchy" in reality, even if the RM isn't using graphDB technology to manage them.

There are two *mechanisms* of state transition (not reason, but mechanism) we might tend to see: either automatic or manual. Automatic means that you don't need consent or any other active participation by parties for the transition to occur. It times in (little Johnny grows up), or a law now goes into effect, or a sensor reports a high enough temperature, or whatever. You might, however, need notification and other workflows. Manual is when parties do need to take action, such as consent or permission or setup (such as registration). Applications could need custom workflows in either case.

Next steps: Apply the three FB use cases and Nancy's healthcare use cases to UMA and our state machine and try to make them "work" visually and in writing.

We don't have a formal meeting next week. Cigdem will inform whether she can hold an ad hoc meeting at either our regular time or another time. We do have a meeting on **Aug 20**.

2019-07-30

Attending: Eve, Andi, Cigdem, Nancy, Vladimir Prihodko (works at Lush Group, has implemented UMA1 there), Domenico, Colin, Tim; regrets from Lisa

We examined the outputs of the ad hoc meeting of last week (in slides form), comparing to our use cases spreadsheet (will grant access on request). The use cases drive the state machine. So you enter it at whatever point depending on the real-world circumstances. There's a question about which system you're interacting with in each use case. E.g., an employee system is an RS that holds hours worked, paid time off, etc., and personal data management is divided up between RS's. We may have to try this with use cases from different industries and see how it comes out. But presumably any single "run" of the UMA protocol only involves a single cohesive system that interoperates.

If you're operating at the technical layer, by definition it sort of all "works out" because the protocol defines how things work, but at the policy layer, questions creep in: having multiple ROs (RRAs) is like science fiction (the multiple-worlds hypothesis 🙂). If multiple RRAs can set policy, does one's policy override another's, or do they combine, or do they get "shares" in the policy in aggregate, or what? These could all be different use cases. We know of some common ones, like where the parents (or set of guardians) for a child have to all agree to release the medical records of a child. But we don't have to figure out how to implement this, just acknowledge that ASOs need to be free to meet the needs of their users.

We could leverage a table as well for describing the states.

	DS involved	DS not involved
Single administrator	State1: Self-management	State3: Management under care
Multiple administrators	State2: Co-management	State4: Co-management under care

That might be a useful way to discuss those two factors in common for a lot of the scenarios. We're not sure if there is a third dimension. Could any one resource have RRAs that have to "split" along these lines? E.g., **Joe** and **Jane** have a joint bank account. **Joe** is competent to self-manage it, but **Jane** becomes ill and wants to designate a third party, not **Joe** but her attorney **Jim**, to manage the account for her. And *then* something happens to **Jim**... It's delegation turtles all the way down. How realistic are some of these? There are legal constraints on continuing to add further and further delegates. But at some level, US laws such as RUFADAA (which is more widely adopted than ever – later Tim provided adoption status in email) could provide a basis for at least the first layer – or maybe as far as two layers? – of delegation to be handled at a "policy layer expressed by technical means" with UMA's help.

Is "policy layer" the right name? The overall generic name has been "business-legal" but that may be too awkward. Let's continue this topic.

Would it be possible to define the policies themselves to make them interoperable? To be discussed.

Questions we still need to discuss as outlined in the slides:

For the care scenarios, do we need a primary carer/super admin type of role or are all entities in RRA equal?

Federated authorization or not? Arguably, personas of a DS and a single DS have the same identity verification risk. If personas, then "DS" in the state machine is replaced with "persona". The key assumption: KYC works – if that fails, all the rest fails.

2019-07-16

Attending: Eve, Lisa, Domenico, Cigdem, Mark, Tim (regrets: Andi, Colin)

FYI, Eve can't make IETF 105 in person after all. She may try to attend the OAuth.XYZ portion of it remotely (it's being presented on Tuesday next week).

The chosen term Representative is now in the Legal Parties tab.

AI: Tim: Develop a suitable definition for Representative, sourced to any legal sources as he sees fit (akin to how other definitions are done in the original report).

What do we say about consent? We bundle a series of actions that the RRA is "authorized to delegate" to the ASO, including "access control, consent, and licensing functions". What does our repeated phrase "manages the sharing of X's resources" in our use cases? With respect to UMA, it specifically means the actions of the AS. Thinking about something like a "vault" or "wallet", UMA doesn't have a technical entity like this, though we know of at least one extension that does, so maybe this could come into play officially eventually.

What is the "state" language about? It's techie language. Many of the use cases would reflect life cycle changes, as in literally a human's life cycle (connected to parties). These tend to be relatively slow and predictable in the scheme of things. Permissions might need to change in response to this. Some might be much faster and more dynamically changing, like temperature changes or associations between people that are more short-lived, like Uber rides or similar.

Why is UMA itself insufficient for this relationship work? UMA-based sharing manages only what a resource owner can control. Adding the "IRM" layer enables us to capture both all the steady states where the "endpoints" aren't the ultimate "end parties", and all the transitions between steady states. So our next significant work here is to define this state machine.

AI: Lisa and Cigdem: Press ahead on the state machine depiction approach.

A "task force" will work on this and we will next meet two weeks from now.

2019-07-09

Attending: Eve, Lisa, Adrian, Domenico, Andi, Thomas, Tim, Colin, Nancy (regrets: Cigdem, Nancy)

Eve briefly walked through the paper that she and Lisa have submitted to the IEEE ComSoc CfP. It's called Beyond Consent: A Right-to-Use Licensing Agreement for Mutual Agency. The argument made by the paper is that digital consent and Terms acceptance (perceived as consent) are failing and don't meet a strict definition of consent (Nancy Kim's *Consentability* framework is used), and using a Me2B lens (centering on the user of the digital services) shows that a licensing agreement is more appropriate. A taxonomy for license agreement contents is proposed, and some challenges are discussed. The paper points to the UMA report where a license is already proposed, but starts "earlier" in the personal data usage chain to be more comprehensive.

What about consent receipts? They record the results of consent. Lisa and Eve meant to cite them in the paper and hope to have a chance to add this.

Adrian discusses a link between DIDs and standards such as UMA; the link is what's called the service endpoint. The DTD standards don't talk about what the service endpoint might be. They're trying to put a personal data store there. Thomas spoke with Microsoft's Ankur Patel (who was with Preeti Rastogi) at Identiverse about the challenge of getting personal data into wallets/personal data stores. There seems to be a lack of recognition of this challenge.

It should be kept in mind that the paper treats personal data permission use cases that go beyond UMA.

Tim recommends removing the "IANAL" disclaimer in the paper! Lisa and Eve have both worked with many lawyers and have sourced ideas from legal experts.

We started to walk through the Right-to-Use License Agreement (Figure 5 in the paper) and analyze which would already be baked into some artifact, such as the RPT, and which might need to be captured separately. For example, the digital asset, grantee (licensee), and actions, would be captured as resource ID, requesting party and client, and scopes. But other information might need to be captured in some other structure, with perhaps a link off to it that is stored in the token. If a requesting party or client received and agreed to such a license, maybe by signing it, the result might be a "receipt". Trunomi has something called a "certificate", which sounds similar, as their consent receipt.

Eve will email a copy of the paper to all interested to review.

2019-06-18

Attending: Eve, Lisa, Thomas

Lisa has uncovered some discomfort with UCITA in communities we care about, which may give us pause about our references to it in the Business Model report as a source of uniform language. See this source for some of the controversy. Section 2 of our report says "A possible challenge to implementing a user-centric access sharing protocol has been the lack of a set of uniform default contractual rules for the exchange of personal digital assets. Fortunately, UMA may leverage the Uniform Computer Information Transaction Act ("UCITA") as one source of default contractual rules upon which the licensing of access rights to personal digital assets may be based." We haven't yet canonicalized *any* standard boilerplate, so we're probably not in any danger of "exciting those antibodies", but it's something to watch out for. Lisa notes that, from the P7012 perspective, fast progress is desired – but when the totality of options is presented, people can easily get overwhelmed.

Is it practical to define a Data Subject as either a natural or a legal person, as has been suggested? Architecturally, yes, it could be. But our analysis suggests that defining it this way is unhelpful and possibly harmful, because:

- UMA's primary aim is to aid "Alice" the individual (a natural person)
- GDPR and many other laws/regulations/policies define data subjects as natural persons and exclude legal persons from this role
- Conflating natural and legal persons is generally confusing

So let's stick to **Data Subject as just a Natural Person**.

We've been looking for a term for the person (Person? Natural/Legal? or only stick to Natural?) who is in a "proxy" role in our business use cases. If Proxy doesn't work, Tim was suggesting Representative or Legal Representative. Examples he has provided: "Personal Representatives, Executors, parents of minors, guardians appointed for minors, Conservators, guardians for elders, corporate proxies, etc." While Data Subject is human, the Representative could be a Natural or Legal Person. Given this, that's probably a good rationale for *not* adding the word Legal on the front of the term, since that makes it ambiguous (iow, it's a "legal representative" whether it's a legal person or a natural person). So let's call it **Representative**.

The point of finding and defining this name is that it's a name for a the party when they're acting on behalf of the Data Subject, not acting in an UMA flow capacity – even though (Representative == Resource Rights Administrator) in the same way that (Service Provider == Relying Party) in federated identity. The first role is about inherent value-add and the second role is about specialty protocol dance.

In the future, we might want to have a term with a definition like this, but we don't yet:

(some phrase): *The Legal Person to which a Protected Resource relates.*

We know there are use cases like this – "enterprise UMA" use cases. Maybe we just say "Legal Person" for now, and we talk about the Protected Resources that relate to them; those resources are not personal data of the Legal Persons (though the resources may contain the personal data of individuals).

2019-05-28

Attending: Eve, Cigdem, Domenico, Lisa, Colin, Adrian, Tim

Suggested agenda:

- Capture additional use cases people have on their minds (up to some limit so we can get to the other items too)
- Decide what to do about "Proxy" and "Org equivalent of a Data Subject" as legal roles
- Start to figure out/capture technical "links-to" relationships that enable us to go from "code" to "prose"

There is work going on on a protocol called ALIAS that is built on OAuth/inspired by UMA. Mehdi Medjaoui is involved; Eve saw a demo. She's begun discussing with him and his colleague the potential for getting together on something like "privacy-enhanced UMA" here based on the fact that our business model work is going in this direction (sewing together the technology with the legal and business layers) and folks like Identos have been extending UMA explicitly in a similar direction. ALIAS has an artifact called a "bind token" that cryptographically binds the RRA (ro), RSO (rs), and ASO (as) (their version). Did Airside Mobile do something like that too?

Adrian points out that "business model" in a business context usually means "How does this organization make money?" Other people also thought that was the meaning. That's not what we've been meaning by it – rather, we've meant the set of relationships that obtain among the parties (vs. the relationships that obtain among the technical entities as defined by the specifications). So, should we call this a "legal framework" instead? Business-legal < *something*>? Something very important hinges on the AS-RS separation, which conveys a main benefit. Let's try **business-legal framework** for now and see how it feels. Is **relationships** a word we can work with somehow? Domenico suggests **business relationships for UMA deployment**. Sometimes Eve has used "deployment use cases" to reflect the fact that they have concrete jurisdictions, people, companies, etc.

Do we actually need a word for "Proxy"? Our diagrams and text use cases in the original slide deck don't seem to need it. Tim hasn't found a single word that perfectly fits the concept, and using Proxy may not even be legally accurate. Note that this sort of (Agency contract, Access contract) delegation is a delegation of management, not of ultimate sharing. It's a kind of "delegated administration for consumers". But certainly, those liable for releasing access will want proof that the licensor has delegated authority.

As few words as possible between the original concept and the UMA mapping, the better! But do we still need some consistent word for the person who is the non-data subject administrator, perhaps Representative?

Perhaps we simply need to recognize the patterns in play by adding their relevant relationships (e.g., the addition of the relevant parties, relationships, and legal devices). This could work elegantly. Eve will try to convert over to that. (Cigdem has also sent thoughts to Eve.)

Adrian wonders if we have now collected enough use cases, since we've now sussed out whether the AS or the RS is the focus of the use case. Eve's theory is that, while there is likely a small and finite set of patterns, it's still interesting to capture further use cases as they arise, since real life is "messy" and we may learn from corner cases.

2019-05-21

Attending: Eve, Nancy, Lisa, Tim, Cigdem, Adrian, Domenico, Mark

Our goals in filling out the new use cases spreadsheet are:

- Make crystal-clear the legal-technical mappings
- Link to definitive definitions (and figure out which concepts need new terms and better definitions, and figure out if we need to refine our definitions in a new report vs. the first report)
    - E.g., define a new Proxy/Agent/whatever term? Define a term for an organization that is the equivalent of a "data subject" (the information is "about it") only it's not an individual?
- Get people to put their hands up in identifying which use cases they care about, and add new use cases
- Figure out how to incorporate all the necessary complex information into licenses, such as covering Requesting Agents and Requesting Parties as required (Dr. Bob and the institution he works with, e.g.)

Our hope is that there are some kinds of licensing that you can boilerplate, a la Creative Commons. That has been the premise behind our looking at CommonAccord.org, which literally has a GitHub system for reuse of legal text (and the Ricardian system in general: prose, code, parameters). Also, jurisdictions will suggest some similarities. (See our very early UX work.) But there is extreme variability that can arise based on things like the types of resources and scopes; for example, "getting some data" is different from "controlling a smart camera".

Eve's "new permission taxonomy" lists five possible axes of control: Scope, Grantee, Environment, Usage, and Downstream. UMA as a technical layer enables three of them. Legal licensing is needed for Usage (e.g., preventing usage for marketing purposes) and Downstream (e.g., specifically, preventing sharing with a further requesting party that doesn't share the same AS as the initial requesting party).

Adrian suggests following up to learn about the work of this group.

2019-05-07

Attending: Eve, Adrian, Cigdem, Domenico, Colin, Lisa

Eve proposes a use case around "relationship-driven policies" that themselves drive UMA authorization assessment and RPT issuance. The mapping to artifacts goes back to a PAT that the RS, as an OAuth client of the AS, has been issued due to the RO's approval.

So far, our nascent second report has an abstract of *"The premise of the draft report A Proposed Licensing Model for User-Managed Access is that UMA enables the individual to centrally manage access and use rights with respect to personal digital assets by converting permission tokens into machine-readable licenses. This companion paper outlines a formal means by which this conversion can take place."*

- Does this work? Is the language apt? The IEEE 7012 work is relevant here.
- OAuth talks about access tokens and we have used "permission tokens" in order to be relevant to lawyers. Are they connected? How big is the gap, if any?
- What does "machine-readable license" mean? It wasn't intended to be a straight sub for "smart contract", but rather something like a dispositive connection to a permissioning technology. Is a consent receipt a machine-readable license? It's meant to be in arrears, and maybe it could memorialize a license. Should there be an IEEE/Kantara liaison? There's some overlap and some intended informal liaising for such purposes.

What is the ideal outcome of the second report? There are real-life human (relationship) changes. They map into something legal, and those map into something technical. If we can identify the latter two, we can help a variety of parties create (interoperable) solutions in order to reduce friction in solving problems of convenience, trust, compliance... **We need to crisp up the value-add of our deliverable next time.**

How do we bite off only enough that an RS would be willing to do it? Since enough power licensed by the ASO to the RSO ("licenses-perm-granting-to" and a lot of downstream stuff) hinges on the PAT, and that hinges on a client agreement, the strength of that agreement – and perhaps any technical stuff that can strengthen it such as signing – becomes important. The other half of what's interesting is all the requesting-side mappings, including the "licenses-perm-getting-to" (bifurcated into CO and RqA). We have question marks there around the fact that the front channel is used a lot for the RqA, and that's invisible to auditing a lot of the time.

2018-03-02

Attending: Eve, Colin, John, Kathleen, Mark, Tim, Thomas, Bjorn

Status update on the publication of the business model doc: It is finally posted! The listing on the Reports & Recommendations page needs to be fixed, but you can point people directly here.

We kind of need a UML diagram (or some kind of graph that lets us label the arcs) to express our delegation and licensing relationships more clearly, and we may need more role names for clarity as well. For example, a Resource Owner "is-a" Data Subject Representative (or whatever we want to call it – that's what GDPR calls it) but we have gone directly from DS to RO in our delegation mapping step. It would be clearer if we had a formal model that went:

- DS delegates-X-rights-to DSR (the mechanism for this might be law, by contract – e.g. a will...)
- DSR is-a RO
- RO delegates-X-rights-to ASO
- etc.

Can we get hold of a UML modeling expert to ensure we don't miss any of the relationships and roles?

In terms of finding one – or multiple – use cases to map onto the model, Eve is talking to 2-3 candidates about this.

The complete list of agreements that seem to be possible in our model so far (the purple ones weren't mentioned in the business model paper yet):

- Agreement that turns a service provider into an RSO (wasn't included in business model report)
- Agreement that turns a service (or app) provider into a CO (wasn't included in business model report)
- Agreement that enables a Person to act on behalf of a Data Subject [which puts them into position to act as a Resource Owner -- otherwise RO=DS]
- Agreement(s) that delegates authorization for an ASO to grant access permissions on behalf of an RO (typically Ts & Cs, privacy notice, EULA...)
- Agreement(s) that delegates authorization for an RSO to manage resources on behalf of an RO (typically Ts & Cs, privacy notice, EULA...)
- Agreement that enables a Person to act on behalf of a Requesting Party [which puts them into position to act as a Requesting Party Agent -- otherwise RqPA=RqP]
- Agreement that delegates access seeking to a CO on behalf of a Requesting Party
- Agreement that delegates permission to know and persist personal data to an ASO on behalf of a Requesting Party

Jim H has started on a CmA version of the model.

It sounds like we need to do this for the POC:

1. Complete the formal model (and likely express it in a formal way)
2. Construct the set of agreements and licenses (or whatever the latter end up being in the case of individual permissions?) in a skeletal CmA format
3. Use the CmA format to invite UMA deployers to work with us on testing the format by applying their use cases

A central element of our proposal is that licenses have the right design characteristics to give ROs (acting on behalf of whoever the DS is) to enable the autonomy, reciprocity, and objectivity needed above and beyond contracts. Regs variously require policy elements of the contracts to enable the individual to have certain rights. "Right of action" in healthcare is one.

2018-02-16

Attending: Eve, Colin, Jim, Thomas, Tim, Sal

The [International Association for Contract and Commercial Management (IACCM)](#) has 40K or so members. It's old and venerable and works on "the contract problem". For large businesses, this looks like: coming up with a legal vs. a business deal, being unable to react to events, having insufficiently agile contracts, etc. Marketing, purchasing, and other business roles are engaging. They have chapter and annual meetings: Europe, NA, and Australasia. Jim is now head of the tech community, a new role. His goal is to network and bring together the threads around improving transacting. Tim Cummins founded IACCM. They are increasingly aware that businesses must play in networks and supply chains. Does a KI/IACCM liaison make sense? Colin and Jim can talk offline about this possibility.

What is the premise of machine readability in our model? The terms of the license in our model could be conveyed through the form of a smart contract or in some machine readable form.

Jim: Smart contracts aren't really contracts. Thomas: All smart-contracts need (a) Digital Identities and (b) some way to control ledger access by those identitiers. I would think Kantara is best position to address the "edges" of a blockchain system. Colin: Agree with you Thomas. Actually the folks starting to operationalize ID on the BC are opening up to conversations about how to assure those pieces.. after having said initially 'we're new, we're hip, we know everything'.. a year on there is some appreciation that the essence of many aspects of the identity assurance puzzle don't change - only the environment that they operate in.. Jim: And (following Thomas) the blockchains may be something like the "edges" between resource and authorization nodes.

Regulators are only nibbling at the edges of really fundamental and core identity issues, which KI is perfectly positioned to weigh in on. And UMA2 is well positioned to answer permissioning questions.

Are all the smart contract questions too big for now? What if we just go for the mappings that enable machine readability if you want it? If you can end up with stable text that can live at the end of a URL, that may be sufficient for now so that we don't go nuts trying to solve the world's smart contract identity problems before testing the UMA business model proposition in front of us.

Our business model report focuses on the resource owner/data subject and requesting party/requesting party agent needs first. But what about the IACCM "personas"? For example, slide 11 in the Legal Role Definitions deck shows "delegation relationships" we haven't yet captured in the report: how a service provider becomes an RSO or a CO. They would get "purple arrows" because they're partially based on UMA technical artifacts (namely, OAuth client credentials).

Would we want to output clause templates? Sounds like it. Jim: Think in terms of general engagement, specific engagement, deployment. See [this example](#) he's working with. The use case: A company is buying data processing or consulting services. In this case it's US domestic, but we want to make ours GDPR-enabled so that it compels a particular pipeline of enforcement actions. With all the subcontracting, Jim guesses that GDPR will dictate a lot of the contractual terms.

We have to pick a use case. Two candidate we've talked about include:

- Origo/Pensions Dashboard
- Health IoT

For our use case, will it require multiple natural languages? We'll have to see.

It would an uninteresting test of UMA not to include the Federated Authorization portion of the specs, so let's assume we're including the "Licensing access granting permissions on RO's behalf" relationship portion.

We also have to enumerate the legal devices that would be created. There are static and dynamic ones.

- Agreement that turns a service provider into an RSO (wasn't included in business model report)
- Agreement that turns a service (or app) provider into a CO (wasn't included in business model report)
- Agreement that enables a Person to act on behalf of a Data Subject
- Agreement(s) that delegates authorization for an ASO to grant access permissions on behalf of an RO (typically Ts & Cs, privacy notice, EULA...)
- Agreement(s) that delegates authorization for an RSO to manage resources on behalf of an RO (typically Ts & Cs, privacy notice, EULA...)
- Does the PAT simply link to all these previous agreements in order to establish that the RO has agreed to the "licensing of access granting permissions on RO's behalf"? We think so
    - It's possible to profile UMA to require that the basis for PAT issuance is interactive vs. silent – this is something we could consider building in to the agreements above, to ensure that a human RO is given the change to consent in a GDPR-compliant way
    - There should be a way to force interactive user consent again if there is a new version of the agreement available
- (to be continued next time - licensing of permissions and the requesting party side of the equation)

Note that we have no meeting next week, Feb 23.

2018-02-09

Attending: Eve, Tim, Bjorn, Kathleen

We don't yet have Tim in a WG [leadership team](#) role with a title! And he really deserves one. 🙂 How about something like "legal adviser"? Let's bring that up at the next whole-of-WG call.

AI: Eve: Bring up a motion on the next WG call about Tim's leadership team role.

We made further tweaks to the [2018 charter refresh proposal](#) to reflect the Legal work stream.

There's general agreement that the Consent Receipts format, so far, accounts for general opt-in consent as known and used today, but it doesn't account for the kind of scope-grained, asynchronous consent/authorization/policy setting and withdrawal that UMA enables.

Things to consider in our business model: Can the ASO be a true Agent even in the use case where the ASO is, say, your IdP and wants to be your trusted AS, but doesn't hold any of your personal data? All your protected resources are held in third-party RS's, so the AS hooks up with them through PATs (OAuth) in an overt way. The challenges would be that the ASO can still learn about:

- Which RS's Alice uses (Schwab vs. Fidelity) – could be tempted to sell this information to advertisers

- Some notion of the nature of the protected resources, through the metadata uploaded as part of resource registration (e.g., see the HEART profiles, which point to FHIR Resource types) – would know what resource is an EHR – could be tempted to sell this information
- The requesting parties associated with Alice and information about them (graphs of relationships), to the extent of whatever claims were collected – if the ASO isn't Facebook or Facebook-based claims aren't collected, this presumably limits the scope of discovery of who everyone is, but the risk is there

The theory is that these risks could be managed much better in the case of a "personal authorization server" (which is what Adrian and Michael Chen have built into their HIE of One implementation) because the ASO's business model would not have to compromise for its *single* resource owner. Perhaps it's as simple as the "pie chart" showing the ratio of money-to-data-to-attention that the resource owner pays the service. If a SaaS company offers an AS to millions of potential ROs, and the business model is such that the ROs don't have to pay any money outright, then it's really hard for the ASO to be a true Agent. In other scenarios where the AS and the RS (or one of the RS's) are colocated, then the ASO is already "compromised" by this:

- The ASO-that-is-also-the-RSO holds some/all of Alice's resources and thus can see that data (so "trust mitigations of trust attacks" – auditability of these relationships – may be needed)

Further, if the ASO also runs the client (or one of the clients) that the RqPs use, then the ASO is already "compromised" by this:

- The ASO-that-is-also-the-CO is gaining access to resources using tokens it issued itself (so "trust mitigations of trust attacks" – auditability of these relationships – may be needed)
    - Kathleen notes that they have been discussing signing of contracts

Tim can join the WG calls as needed for when we combine this work with overtly technical topics. Kathleen (and Mohammad) generally can't join the Thursday calls, so we can either schedule ad hocs as necessary or perhaps use the Legal call time as necessary.

(A reminder: Our "Legal role definitions" deck is here. The paper will be published officially by next week, but you can also find the unofficial version, which has all the diagrams, in our mail archives here.)

AI: Eve: Ask colleagues if their health IoT storyboard could be used as a starting point for the POC idea.

AI: Eve: Reach out to Jim Hazard to alert him to the POC idea.

2018-01-19

Attending: Eve, Colin, Tim, Adrian, Mark, Kathleen

Agenda:

- Final edits to the business model doc draft 7c, taking into consideration comments collected during UMA telecon 2018-01-18
- Get the draft ready to send to WG members for review in PDF form so we can do an e-ballot for draft WG report publication

Regarding the Vermont legislation: Though the overt focus is blockchain, Colin has some thoughts about the notion of a "personal identity trust company". Can we simply have a conversation with the legislators? Yes! Tim reads the bill as being about identity brokering. Which others could also provide input? Perhaps those with an eye on the legal landscape as well as identity and privacy. (John W?) The bill authors have expertise in banking law. Adrian could talk to them as well – turns out he's already aware of the effort. We don't want to necessarily derail their "blockchain" focus, but we want to introduce them to our notion of the ASO as an agent and the power of that. Colin can help with those experienced on the brokering side, and mentions some comparative country ID efforts. Adrian mentions India's identity system Aadhar and their efforts towards greater privacy; he's involved on the health records side.

AI: Tim: Make introductions as appropriate between the Vermont legislators and the UMAnitarians.

Business model paper:

Elevator pitch in doc:

A Proposed Licensing Model for User-Managed Access

Can we describe our audience much more directly?

Old: *This paper is intended for audiences who are business-knowledgeable and experienced with legal devices.*

New: *This paper is intended for professionals in the areas of law, privacy, risk, compliance, security policy, and business policy, particularly those responsible for building and running UMA-enabled services.*

Latest subtitle/elevator pitch: *How the UMA access sharing protocol together with a licensing model for personal digital assets enables user-centric control in the network-based information society*

Mike's, with additions: *How UMA enables an individual to control access to their personal digital assets in the information society <<when UMA-enabled services are mapped to a licensing model>>*

Playing around/discussion: *UMA protocol (associate it with "privacy" and "privacy instruments"?) and licensing model*

Further work: *How licensing enables personal digital asset access control for individuals using services enabled by the UMA protocol, heightening privacy protection <<and compliance?>>*

Arrgh, so close! Tim and Eve will try and wrap up all the remaining comments in the doc by Monday and get the e-ballot out.

2018-01-12

Attending: Eve, Colin, Tim

We worked through paper 7c live on the call. Eve finished her action item to add a scope/audience paragraph live, too.

There are Code of Conduct materials extant, both external to Kantara and within Kantara (including pointing to those other sources), which can be built on to bootstrap an UMA-related "company code of conduct" assessment offering (and even a similar OAuth-related substrate). Some discussion went on about this in the past week. The current Business Model paper, and probably a series of other short papers in some "call tree" order (paper A cross-refers to paper B etc.), would be formally referred to by some set of tools. If these were technical specs, those linkages would be obvious – so let's be sure to make them obvious. Paper draft 7c is getting ready to be a "beta" version for external review. The WG should vote it upward as a Draft Report if so, and we'd have to see if the LC then needs to review it, but then it could potentially be put under the Reports & Recommendations page as a PDF download.

The legislature in Vermont is developing a law that has UMA-ish elements. They have the concept of an "autonomous agent", which sounds like an ASO! Tim mentioned that UMA has solved this issue. They may need someone to testify on such matters.

AI: Tim and Eve: Work through the paper on an ad hoc basis early next week.

2018-01-05

Attending: Eve, Tim, John, Colin

Agenda:

- Get the Business Model paper ready for its closeup (International Privacy Summit presentations on Jan 29)
- Discuss satellite papers

Tim has prepared an "experimental brach" of the business model paper, as a draft 7b. It plays off of the semweb/ontology topic. We need to briefly state some things up top (maybe a last paragraph of the Executive Summary?):

- The audience?
- What the UMA protocol accomplishes (the parties by name and abbreviation) – possibly then pointing to a separate paper
- What this paper accomplishes (UMA-compliant licensing)
- A very short example
- The scope (UMA2, not UMA1, and Alice sharing, not AliceCo sharing)

Should we entirely avoid the abbreviations for the parties? We avoided ever using the abbreviations in the specs themselves. Eve believes strongly that we should include the diagrams; we could potentially spell out the terms rather than using the abbreviations.

Is a cloud storage service provider considered a custodian? a bailee? Is bailment transitive? The "Challenge of Personal Ownership" section currently mentions this question. The bank/lockbox bailment analogy is often made outside the digital resources context. These roles should, the theory goes, love UMA. John suggests distinguishing "records", "information", and "data". Owning a record shouldn't mean owning the data. A service provider who's a data controller for an individual doesn't own that person's data. Information is the stuff described by the (digital) data. The subject has rights over the data. When Alice has concerns over her privacy, it's about the data. To date, the paper hasn't dealt with these distinctions. UCITA's definition of "information", unfortunately, is not so nuanced.

If "protocol" is scary, then "standard" would do. If "method" is likewise scary, then "system" would suffice. Talking about "meeting customer expectations" also acknowledges that our audiences are likelier "B2C" than individual data subjects. The first time "standard" is used, maybe it could be prefaced with something like "machine-to-machine communications standard" or "digital communications standard" or something.

In the world of legal devices, contracts (today, anyway) are the "design-time" devices that are stable and heavier relationships, and licenses are lighter-weight relationships that can be determined at run-time. UMA does not require any of the devices to be signed, but in combination with receipts, they could be. Any artifact passed between two entity, signed or not, is visible as identical to both entities and therefore auditable.

Let's continue with draft 7b.

John hypothesizes that, once data is shared all over the place under UMA, a lawsuit will be filed under UMA licensing. Eve naively assumes the judge should be able to inspect the contract and license wording that flows upwards from any obvious template language we provide. Would that be so? Maybe class actions might not be so simple!

AI: Eve: Write the short scope/audience paragraph.

AI: Eve: Try to turn the old Legal Considerations doc into a separate paper we can call out to that provides (hopefully our new set of) UMA use cases and an extended explanation for a business/legal audience of the technical side of UMA.

2017-12-22

Attending: Eve, Kathleen, Tim, Bjorn

We now have a draft 6 of the paper thanks to Tim's efforts! He considers it nearly complete, which is amazing! He has resolved all the must-do comments. There are other "philosophical discussion" comments that still appear in the draft 5 GDoc, so we haven't lost any of that, but we need not address them them all.

AI: Eve: Create GitHub issues with the "trust" label (see others here) for any "real" remaining GDoc comments.

Eve talked to Origo and they're interested to check out the UMA Legal work. See their Aggregating and Sharing Pension Information case study video here.

There's a Kantara-supported International Privacy Summit on Jan 29 in London. Eve is likely traveling to speak there (if not, then presenting by Skype!). Subject: Something like UMA Legal+Consent Receipts. So let's get the model doc ready for its closeup. Kathleen is suggesting not just throwing consent receipts, but also an acknowledgment, or sending the receipt back with an expiration date, so you don't need a new artifact. She and Mohammad are discussing this. She can share a swimlane.

2017-12-08

Attending: Eve, Devon, Tim, Sal, Jeff

NO MEETING on **December 15** (next week). YES MEETING 🙂 on **December 22** (the week after). That will be our last meeting of the year!

Regarding our planning around getting the framework ready for end-of-year review, let's discuss the definitions. We should put them in the back, in an appendix, because they don't flow narratively. We should cross-refer to them from the narrative text and we could support that reference with some illustrative diagrams to help make sense out of them.

The legal definitions have an RSO definition that is nearly there, but seems to be missing the connection with the RO. E.g., if you look at the definition of an RS in the Grant spec, it is "*A server that hosts resources on a resource owner's behalf* and is capable of accepting and responding to requests for protected resources."

The PCT is not like an RPT; it doesn't mean "I've got access to all this stuff, so please let me in to it again", it means "I've got Bob with me again, so can I please get access to whatever stuff Bob gets access to?"

Though the UMA specs avoid the word "consent" (much as the OAuth spec does except in a single, probably accidental case), this framework doc need not, and can actually advocate for the case that UMA is directly and concretely relevant to data subject consent. This is because we can make an *unbroken* chain from technical artifacts and entities to parties' delegation of authorization and permissions. Or, depending on the audience, the reverse! Do we really need an unbroken chain, or is that technical/spec thinking that isn't applicable in the legal-framework case? It's Eve's theory because of the whole "prose objects" approach, "smart contracts", and the like. We'll have to see. (Tim suggested Devon could put some thinking into that. 🙂)

There is a specific use case when Alice herself acts as the RqP. In this case, the CO role is outsized because all the weight of "who Alice is sharing with" gets placed on who/what the CO is.

Could the client (or Bob) require that a consent receipt be issued around the RqP-CO delegation relationship? This could be considered a best practice, as part of the UMA-CR joint work.

AI: Tim: Revise legal definitions before inserting into the doc, to ensure they align with our latest analysis of the delegation and licensing mechanism and that the RSO definition connects with the RO in some fashion.

2017-12-01

Attending: Eve, Kathleen, Tim, Devon, Theresa, John, Bjorn, Mark

Agenda:

- Review what we learned in yesterday's joint UMA/Consent Receipts call
- Look at new language in Business Model Paper 5
- Figure out what we can accomplish by the end of the year

In HIPAA, there's "consent directives". "Release of information", "patient right of access", and "authorization" all mean the same thing. There is no authorization required for TPO (treatment, payment, and operations). For operations, an entity has to have, or have had, some kind of relationship with the patient. A "covered entity" (a term of art) would then have "HIPAA consent" (different from regular consent, which the patient has no ability to *not* grant). There's implied and specific consent. ISO standard 17975, **Health informatics -- Principles and data requirements for consent in the Collection, Use or Disclosure of personal health information**, lays this out. HITECH has now been absorbed under HIPAA. If you pay for any services completely out of pocket, then the provider must honor your "dissent" (an opt-out). "Grantor's choice" is where the provider has no choice; it's another kind of directive regarding a "right-of-access request". Can we map UMA's sweet and sour spots to these definitions of consent, for various laws and regulations? The paper needs to do this at a shallow level to be credible; the tools need to do this at a deeper level.

The theory is that our framework can lighten the load of anyone building their legal/contractual/licensing layer not only through templating but also, potentially, through business process modeling. This is where a CommonAccord workflow can help as part of the tooling.

AI: Eve: Talk to Andrew about whether our WG can get access to the ISO standard through KI's liaison relationship.

AI: Eve: Work on legal diagrams.

AI: Tim: Add legal definitions to the paper.

2017-11-17

Attending: Eve, Jim, Tim, Kathleen, Colin, Bjorn, Mark

Tim asked: To avoid many many contracts, is it possible to have all parties go through the ASO?

- Eve: Most circumstances today leave Alice accepting the ASO's and RSO's ToS/PN as a contract of adhesion. (Also, presumably, the use of each RS comes at a time when Alice decided to use it, and the choice to sew together each RS with the AS – authorizing PAT issuance -- comes at an appropriate time also, so, the experience is at her discretion.) Our work to add to templated clause text can strengthen the provisions of these contracts even if there are several preceding the license phase.
- Eve: In some deployment ecosystems, the ASO may be singular because other parties are aggregated into it (ASO + singular RSO, ASO + RSO + CO, etc.). In these cases, the process of accepting an initial (non-UMA) ToS/PN and authorizing a later PAT, e.g., could be rolled into one.

- Jim and Kathleen: In some cases, Alice may have power, either because the RO is "AliceCo" (not directly in our framework's scope for now) or because the law gives her more power (e.g. "patient right of access") – effectively, the Alices have been "aggregated". These effectively give strength to the templating we can do.

Jim: Regarding improving the circumstances of the "little" party wrt the "big" party, you can do it through influencing law (out of scope for our group), you can do it through intermediaries, you can do it through model terms/template clauses (the latter now being this group's favored phrase – so this would be our group's role in this ecosystem), and you can do it through adopting processes that enforce industry behavior that manage compliance risk for "big" parties (this is how others could use our deliverables – maybe IACCM, Pan-Canadian Trust Framework effort, other countries, other industries, etc.). We can't make anyone use this framework; rather, we want to make it as easy to use, adopt, adapt, and profile as possible.

These links illustrate templating and ecosystem chains:

http://www.commonaccord.org/index.php?action=source&file=Dx/Acme/09-EU-US-DataTransfer/Acme_UK/0.md

http://www.commonaccord.org/index.php?action=xEdit&file=G/TechContractsCom-ITMA-CmA/Form/0.md#GeneralTerm.Data.Sec

Eve has reached out to the people at Origo regarding their use of UMA (V1, however) and their potential need for ToC/PN that are protective of privacy rights in a standard way.

https://pensionsdashboardproject.uk/saver/about-the-pensions-dashboard/

http://www.origo.com/news/Kenneth_May_Demonstrates_Delegated_Authority_Pensions_Dashboard.aspx

https://www.ipe.com/countries/uk/uk-government-backs-pension-dashboard-project/10021187.article

Let's plan to focus on this as a candidate case study/use case in the document. It has properties of strong authentication, a likely "design pattern" of sharing we'll see again, a government friend for Alice in negotiating the terms ("you and what army?" "this army!"), non-co-located services/parties, mappings to many of the characteristics of the use cases we gathered in deliverable #1, people who profiled UMA for real use already, etc. It only doesn't have UMA2, but we could work around that.

2017-11-10

Attending: Eve, Jeff, Devon, Theresa, Mark, Kathleen, Tim, Ann, John

Doc homework:

- A very early section (the first?) should present the "pain point" by introducing several broad scenarios, including an Alice-to-Alice, Alice-to-Bob, and Alice-to-org, drawing from deliverable #1. It could introduce the language of "resource owner" and "requesting party". We could note that org-to-whoever sharing is out of scope for this exercise (framework). We could have a very high-level version of the x-and-y-axis scenario diagram that just talks about these two roles. and then the version with the three high-level scenarios.
- In NewSec, we want to make the strongest case we can for our chosen legal devices, and ultimately for our biggest target type of toolkit (templates of some sort).
- Later, we can get into the sub-scenarios we have collected, e.g., Alice as a guardian of a data subject too young to consent etc.

The "collaborative diagrams" in the GSlides need more differentiation and "iconification".

"Model clauses" specifically means they need regulatory approval, so how about "template clauses" or even "clause templates" or something? Templates will do for now.

Let's get more specific about pain points.

"Through a combination of strengthening data protection regulations, justified consumer cynicism and savviness about poor security and AdTech/MarTech ecosystems, and good rationales for data sharing, particularly in the cases of healthcare and the Internet of Things, we're seeing people start to be given just a little more transparency into and control of their personal data. Organizations have never had more incentives to make changes and reduce friction..."

The healthcare construct of a "consent directive" can be directly and favorably compared to ToS opt-in (or, for that matter, opt-out – soon to be made effectively illegal by GDPR) as a mechanism for inviting individuals to express their data sharing preferences in ways that are not influenced by outside actors. UMA enables this construct to be digitized in a standard and repeatable way. This framework enables it to be

AIs:

- Eve: Create two new scenario diagrams ready to put into the GDoc:
    - Very high-level diagram introducing "RO" and "RqP" language
    - Fixed three-scenarios diagram
- Tim: Flesh out the licensing framework itself
    - Possibly this includes the rationale as started in the comment on NewSec

Be sure to see all the new comments in the doc.

2017-11-03

Attending: Eve, Sal, John, Colin, Mark

Eve has proposed a joint meeting at the end of November or early December on a Thursday morning US time to take advantage of existing CR/UMA WG meeting times. The UMA WG is meeting Nov 16 to progress its Draft Recommendations. Nov 23 is US Thanksgiving. Could Nov 30 work? It seems so. Let's try and work on an agenda now.

- Who should attend in addition? Be sure we invite them.
- What are the business/legal/technical touchpoints? Ultimately develop swimlanes that demonstrate technical CR/UMA interactions.

- How does data protection regulation conformance play a role? Not just jurisdiction by jurisdiction but cross-border and internationally (core /common concepts)?
- Is there an end-to-end Kantara vision of an "Alice to Bob and back again" lifecycle and how do CRs and UMA slot in? Ensure it's not just about Alice.

Eve suggests that the respective WG chairs sharpen up the agenda, help determine and invite additional invitees, and make an event invite happen.

2017-10-27

Attending: Eve, Tim, Kathleen, Devon, Theresa, Colin, Jim, Bjorn

Eve was just attending "EWF USA", where GDPR knowledge is dawning.

We put tons of new notes into the legal model doc and did some editing as well.

2017-10-13

Attending: Eve, John, Devon (guest), Kathleen, Tim, Colin, Bjorn, Thomas

NOTE: No meeting next week (multi-determined: IIW and Eve travel), but we do have a meeting on **Oct 27**.

Would Skype be a good way to coordinate a joint UMA-CIS WG session (or session series)?

See our notes in the new "UMA Business Model Paper 5" GDoc. Let's try and have the paper ready as a solid draft for review by end of Q4.

Here's the latest on the situation with the SCCs.

2017-10-06

Attending: Eve, Kathleen, Colin, Tim, Mark, John, Thomas

Who on this call will be at IIW? John.

We revised the existing train-track diagram in web and OmniGraffle form, and created new ones in the slides. We're not quite done with the RPT diagram, and will also have to create "tear-down" versions. We may also want to create different diagrams for different use cases.

We discussed to what extent terms of service/privacy notices ("ToS/PN" in the diagrams) are negotiable. Assuming they are currently not, thinking about the "cascading OAuth" use case where the main AS wants to either give access to resources that the RO doesn't want them to or vice versa, the RS would have to reserve some resources from central protection – which would be inconvenient for the RO.

We discussed to what extent model clause text should be "mandatory" for the RqP side. We're not sure how our legal tools will be used; it could be that we should make the RO-side clause text "MUST" and the RqP side "MAY", but in practice it will all be discretionary to use. Maybe some parties picking it up for use will want to make it all mandatory. Note that some use cases will need clever parameterizing within the text, such as trying to require a service to allow "right to erasure" but ensuring allowance for the service to meet data retention requirements.

Tim will be able to share the draft framework doc by next week.

2017-09-29

Attending: Eve, Tim, John, Bjorn, Mark

Agenda:

- What work remains for us to have a complete legal framework?
- Getting ready to develop one key set of tools: model clauses / standard contractual clauses (discussed here: http://ec.europa.eu/justice/data-protection/international-transfers/transfer/index_en.htm)
    - Could we actually get ours preapproved?
    - Note controller-to-controller transfer vs. controller-to-processor transfer options, something we've discussed in the past

Jim pointed out his "Accordified" GDPR text efforts:

- DE/FR/NL/EN. (Spanish and Greek also available): http://www.commonaccord.org/index.php?action=list&file=Dx/Acme/09-EU-US-DataTransfer/
    - E.g., English language version under hypothetical of a UK sub of a US company transfering data to US parent: http://www.commonaccord.org/index.php?action=xEdit&file=Dx/Acme/09-EU-US-DataTransfer/Acme_UK/0.md
- Source: http://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex:32010D0087

What's the difference between BCRs (binding corporate rules) and model clauses/SCCs? A set of BCRs is a kind of contract that might have SCCs in it. BCRs are developed in an ~18-month process, and have to be approved in all the various jurisdictions where they apply. BCRs have a lot of upfront cost /time, but save the organization cost/time later. SCCs can also be preapproved, but then each contract overall doesn't have preapproval, so you haven't saved as much cost/time upfront. Is it viable for us as a WG or dot-org to seek preapproval of SCCs we write? Maybe we should ask the Art. 29 Working Party about this. This could be valuable for giving UMA/Kantara visibility at that level.

AI: Eve: Discuss the proposition of generic preapproval of standardized SCCs with Colin and Andrew.

[ISO/IEC 29100](#) is the key privacy framework that defines the terms PII subject, PII controller, PII processor (etc.). (We are using Data Subject etc. in our UMA Definitions Annotated deliverable.) In Consent Receipt, they use PII Subject etc. and recommend parameterizing the language accordingly. We can literally parameterize the language in our SCCs to the extent that we capture our language in CommonAccord. GDPR seems like the most obvious target for our completion of the framework, since it's got the right world attention, deadline, etc. Key questions (which we've asked before):

- Can the Data Subject truly serve as their own Data Controller in a permission granting interaction? (Karsten Kinast has said yes before, at least philosophically.) Mark calls this a Master Controller Access Framework.
- If yes, does it make sense for our use cases for [end-to-end licensing relationships](#) to split out to make the requesting party (and client operator?) become either a) another Data Controller, or b) a Data Processor?

Also see the new [slide diagram](#) "Merging RO-RSO, RO-ASO, and RO-RSO-ASO relationship train tracks". This shows how our legal framework would work in terms of contract and licensing clause text dependencies.

AI: Eve: Ask Domenico to take a look at the new diagram.

A key theme of our work is to ensconce the twin to GDPR's (never-talked-about, Article 7) "right to withdraw consent at any time" – the right to consent, or not, at any time (vs. simply opting in or not!). The recitals of GDPR are infused with a data subject's rights but oftentimes regulators can't know what's possible in tech until they've seen it with their own eyes. UMA + consent receipts are powerful in this respect.

AI: Tim: Think about revamping his magnum opus for Eve to "GDoc-ify".

2017-09-08

Attending: Eve, Tim

We agreed that deliverable #3 is complete.

AI: Eve: Let Colin know about deliverable #3.

The [Uniform Fiduciary Access to Digital Assets Act](#) is already in effect in most US states and will require other states to take action quick. The Act calls for the use of an "online tool" – "an electronic service provided by a custodian that allows the user, in an agreement distinct from the terms-of-service agreement between the custodian and user, to provide directions for disclosure or nondisclosure of digital assets to a third person". That sounds exactly like an authorization server! Our legal framework and tools can help with providing that. The Uniform Law Commission has commissioners from every state. It seems they promoted the law pretty well, given the thorough coverage. There's' also the increasingly pressing requirements for PSD2 (January 2018) and GDPR (May 2018). There's also the DIACC Trust Framework effort.

Kantara can take directed funding for efforts to build specific tools, for example contract clauses (tools) for specific purposes. Would that make sense in the case of "directed" sets of contract language for jurisdiction-specific bodies of law covering specific use cases?

Is it time to go to the BL-FIDM list to get attention and review, and see if we can find "customers" for fleshing out the full framework and our first tools in Q4 2017? We think so.

AI: Tim: Draft message to BL-FIDM. Some goals: Ask for review of existing deliverables (some subset?) and slide artifacts (some subset?) to be determined, ask for use cases (and directed funding if needed to hire any further expertise?) for specific tool development, and ask really interested parties to join WG to develop tools actively in Q4 2017 and beyond.

AI: Eve: Cancel next week's Legal call.

AI: Eve: Revitalize our external reviewer list and think about who else to add from IAPP.

2017-08-18

Attending: Eve, Tim, Kathleen, Sal, Colin

Quick discussion related to the CIS WG topic of User Submitted Terms and who is the data controller: John introduces the concept of "Two Solitudes". The more we – in our various WGs and specs! – start figuring out exactly how individuals can influence organizations through these mechanisms, we want to ensure that the "resource regulators" and other influencers get the right idea about how all this applies to the privacy entity roles. The UMA and CIS WGs should plan some joint calls to sync, share wisdom, line up questions, and identify points of interaction and interop between specs and frameworks. Let's plan this at the next Kantara Workshop at CIWUSA17.

Looking at the new version of the role definitions (Word doc on the screen): **UMA Role 4**:

The **Automated Transaction** definition accounts for the fact that these transactions are indeed digital in nature and often automated. So this is an "underpinning" kind of definition. This comes from existing Uniform Law. (Tim's approach is to borrow language from existing laws, so that it's familiar.) He'll add the citations before publication to ensure the text all has the right authority.

The **Protected Resource** definition discusses "All data (and then a long list of data and type of content taken from existing law language)...". Do we mean all data, or specific data/content controlled by an RO, available at an RS, and protected by an AS? "Protected resource" is a specific UMA concept, and this is the sort of "magic triad" that makes it a protected resource. Kathleen mentions the concept of "addressable" as a potentially helpful term to ensure we mean the digital and web-ish kind or even IoT-ish kind. Tim thinks that "information" as defined in a particular body of law he knows (which was it? Fiduciary Access Law? it covers assets after someone has died) could be useful. So do we need to enumerate the kinds of information, or can we reuse someone else's definition? We probably don't need to include a notion of "addressable". Maybe something like **"A digital asset (citation) manageable /managed by a RO"**. Boom! The verb there is just to connect the standard meaning of digital asset to our concept of RO that's in this glossary. Hopefully we don't have to worry about actual explicit liability considerations in choosing this verb yet (though maybe we do?).

We are thinking that the **Authorization Server Operator** definition should say "authorization server" vs "access server", and then cite UMA Grant because it defines "authorization server", and likewise the **Resource Server Operator** definition should say "resource server" vs "host server" and cite UMA Grant. This "gives authority", so to speak, to UMA's official definitions of these technical terms. Likewise in the fullness of time, we'll want to point to the official UMA definitions of the various technical artifacts. (Note that the only artifacts not defined in the UMA Grant spec (or FedAuthz spec - PAT is in there) would be "client identifier" and "client credentials", and those are in the OAuth 2.0 spec: IETF RFC 6749.)

This law Tim wants to use expressly covers what to do about digital assets when somebody has died. UMA has relevance in Digital Death use cases (see past IIW notes for lots of work on this!).

- Even after the RO has died (this body of law is useful even before that time, though!), *as long as the PAT is valid*, and the RPT and permissions are still valid, then an RqP can still get access.
- If the RO does estate planning, then delegation mechanisms could be used at the business logic layer above any one instance of an UMA flow so that a "transfer of protected resource ownership" can be effected.

Though the **Data Subject/Resource Owner** relationship isn't the same as the **Requesting Party/Requesting Party Agent** relationship, we should make sure to have the same kind of "**on behalf of themselves or (the other one)**" construction in the right (other one of the pair).

The Legal Relationships section may have been overcome by events; maybe review it for current accuracy, or remove it. Or possibly we'll simply have the "all-singing all-dancing" paragraph ready for inclusion.

We do have a Legal call next week.

We **DO NOT** have a Legal call on Fri Sep 1. US and Canada people, enjoy the long weekend for Labor Day!

2017-07-28

Attending: Eve, Kathleen, Tim, Bjorn, Mark

(Oops, a bug in the diagram, the sharing scenario labels need to be reversed.)

The nature of the CO's access to the resource is determined by the TOS between the CO and the RqP, particularly in the case of a human RqP (Individual scenario). The client app will have certain access-getting capabilities, and will represent those to their user. Who is paying whom for what? In the case of an Individual RqP using a client app, quite often the individual downloads a free app (possibly with in-app purchases) or pays for an app (this would be for mobile apps, including those associated with smart devices, or even for desktop apps). In the case of a Legal Person using a client app, more often there won't be TOS as such but the client will be developed in-house or done as work for hire by a mobile app dev house.

In both cases, an ASO would tend to be in control of which clients they work with because an AS gets to be "picky" about clients -- clients always bring security risks. But different ecosystem pressures can lead ASO's to be more open or closed. The DirectTrust initial trust bundle (trust framework) was deemed not safe enough so they had to develop a second one. No matter how "dynamic" the process becomes, and the OAuth Dynamic Client Registration protocol is making it quite dynamic, there are reasons for an ASO to potentially want the process to include manual steps, such as lawyers or accountants or other certifiers to get involved. This is precisely where trustmarks and whitelist mechanisms of various sorts are intended to reduce the friction of such manual mechanisms. A client developer could then include a signed assertion in their "SSA" (software statement assertion), which an AS could consume.

The technical artifact we could *eventually potentially* take advantage of here is something in the SSA for dynamic client registration for UMA certification.

There's a nice older diagram, looking like a train track or puzzle pieces, that shows the various TOS's sticking everything together. What might be better is two train tracks, for the RO and for the RqP, where they build up enough relationships with the RSO/ASO and CO/ASO respectively through the legal /technical artifacts and merge. Ultimately we could show them diverging as the licensing relationship gets torn down. (This would be when we know that we've reached the end of designing the legal framework!)

We agree that our goal is to have the legal framework point to the technical artifacts (e.g. as we did in the old Binding Obs), and this is sufficient for now. Various individual tools in the toolkit will ultimately have technical artifacts point to legal artifacts (e.g. URLs containing hashes of the specific versions of text etc.). This latter approach is all implementation-dependent, and so we don't want to optimize prematurely.

We should try a) connecting the definition of Data Subject to Resource Owner, to justify why to include Data Subject, or otherwise drop Data Subject, and b) drop Requesting Party Agent and include "with the legal authority to seek access" to Requesting Party, or not drop RqPA and figure out how to square the circle. 🙂 See the list below, and deliverable #1 for use cases! That rationalizes the two halves of the licensing end-to-end relationship.

The legal versions of the sharing scenarios should actually look like this:

- Individual-to-self (this is potentially interesting legally because Alice tends not to want to impose onerous licensing controls on herself, only on the CO)
- Individual-to-Individual
- Individual-to-Legal Person (maybe this breaks down along the lines of different business models)
- Individual-to- (do we want to consider "Requesting Party Agent" as a role because of use cases?)

WRT Data Subject, we could treat that as a whole separate mapping exercise: DS, DP, DC. Should we do that separately? A topic for another time.

The ONC Trust Exchange Common Agreement work could potentially point to this work to get to a more dynamic model vs paper-based.

AI: Eve: Reach out to Domenico to get his help on creating new diagrams that drill into these detailed relationships.

2017-07-21

Attending: Eve, Kathleen, Tim

We started looking at chart 6 (not sent in email). Tim would like to ensure that we get the pairwise (or more) relationships, legal devices, and artifacts locked down so that he can revise the definitions. The column mentioning "pairwise" should probably take that word out because some of the artifacts definitely represent more than two of the roles.

A *license*, versus a *contract* specifically, gives flexibility around number of parties and promises. It's about tracking permissions. Some permissions are given with some conditions. The license gets revoked if the conditions aren't kept. A license scales better than a contract as a legal mechanism. (Kathleen refers to the licensing model as essentially an authorization model.)

A DURSA (data use reciprocal service agreement), from healthcare, is an example of a surrounding contract that could be used to encompass the kind of license we're talking about. Many of the use cases we've collected – a nanny collecting kids from daycare, giving someone access to a connected car – show that there's a desire to ensure a contractual apportionment of liability across to a requesting party in addition to licensing.

If this were a contract, the RO Alice would be the offeror, the resource access permissions would be the consideration, and the RqP Bob would theoretically be the offeree – but where is the acceptance part? In the technical layer of UMA, Bob might not have acted positively at all to get access. His client software might have "pushed" claims about him without any action on his part, Alice's AS might issue a PCT to his client that represents claims previously collected and the client subsequently "pushes" that PCT back, but that happens silently. Does that constitute acceptance? Do we have to build that into our model clauses? Using a licensing model, RqP Bob *doesn't* have to accept, and any overarching contract layers could build that in only if they want to. Developing contractual-level tools is out of scope for us.

The term "access contract" came from UCITA law. Do we want to use it? It appears in the role definitions, but should we step back from it? Actually, this is between the RO and ASO. So it appears on the wrong line; it should be on the RS-ASO line, not RO-RSO. Or was it correct after all? Are the RSO and ASO both parties to the access contract? The RO becomes *the* licensor, and do the ASO and RSO both become the sub-licensors? It would be helpful, in the "Legal Relationship" column, to always state "*so-and-so* is the *role* for *whom-downstream* of *what* on behalf of *whom-upstream*". E.g., "RSO is sub-licensor on behalf of RO". Maybe we can even construct the columns so that each variable gets filled in, in order, to form a "sentence"?

The RO is at the "head" of the relationships. The RO owns the access rights. (In US healthcare, under federal law, a patient owns right of access.)

The ASO-CO artifact of OAuth client credentials is only related to that pairwise relationship (not to RqP BoB or to the PCT; this should be relegated to a separate row). It could be related to a contract, say, ToS. It's a *potential place* where conditions that are protective of Alice *could* go, e.g. to ensure that the client (later interacting with Bob, e.g. sharing Alice's data, storing Alice's data, etc.) acts in a way that is aligned with Alice's interests.

At the end of the day, the legal aspects have to match the UMA flows and artifacts exactly.

AI: Eve: Try to do an edited matrix in GSlides form for the group.

2017-07-07

Attending: Eve, Kathleen, John, Ann, Mark, Tim

Tim sent definitions according to the "exercise" Eve set. He derived many of them from UCITA, and some from prior UMA materials, including the draft model clauses, and then "protected resource" is new and may need more work.

Let's treat these definitions, not as final model definitions, but our working draft that could be published in deliverable #3 (potentially accompanied by some diagrams showing mapping relationships) to show where the framework is headed. We can all review the document prior to the next call and send comments.

Access Contract: *A contract or agreement to obtain by electronic means access to, or information from, an information processing system of another Person, or the equivalent of such access.* What is the difference between an access contract and an information sharing agreement? The former is a term out of UCITA law, so that's why he grabbed it.

Resource Owner: *A Person with legal authority to grant access rights to Protected Resources; authorized to delegate access control functions to an ASO and to license access and use rights (permissions) relating to Protected Resources; acts as licensor to the Resource Server Operator.* Do both parts of the second clause relate to the ASO, or does licensing access and use rights pertain to the RO alone and not something the ASO mediates? The theory was that the ASO mediates this because it manages and executes/makes decisions on the RO's policies (which it does). Does taking out "*and*" in "*and to license access*" fix this, roughly? It seems so.

Note: In UMA, the policy does not inherently travel with the resource, without some other layer of technology ("sticky policy" technology or similar). Do we need to define **Policy** (or **Authorization Policy**) vs **Business Policy** somehow? Does that add value to what we're doing?

Person: *An Individual or Legal Person.* Great, same as before.

Legal Person: *A legal entity means a corporation, business trust, estate, trust, partnership, limited liability company, association, joint venture, governmental subdivision, instrumentality, or agency, public corporation, or any other legal or commercial entity.* This seems to need just a touch of wordsmithing, e.g., "*A legal entity; a corporation...*".

Protected Resources: *All data, applications, or software in which a Person either has Informational Rights or gives the Person the ability to exercise Informational Rights.* First, should this be singular? Second, regardless of singular or plural, the RSO has inherent authority in various aspects of "what counts as a resource" and "what types of access are possible to perform on a resource", as described here (so possibly it's worth breaking down into **Resource** and then **Protected Resource**, as the former is RS-related and the latter is AS-related?). Is it interesting to consider mentioning "URI" as the location of the resource somehow, or no?

2017-06-30

Attending: Eve, Tim, Kathleen, John W, Bjorn, Colin, Mark L, Sal, Mary

Let's clarify all our role terms and concepts today:

1. **Resource Owner** *(was Grantor)*
   a. **Data Subject** *(was Resource Subject)* (has privacy framework meaning) - could be the Resource Owner or has delegated authority to a separate Resource Owner by law or contract
2. **Authorization Server Operator** - Agent (fiduciary?) of the Resource Owner - role is controlled by the PAT, and in practice it issues the end-to-end licenses because it issues the RPTs
3. **Resource Server Operator** - Data Controller (has privacy framework meaning) - is the Resource Server Operator
4. **Client Operator** - 1st-tier Data Processor (has privacy framework meaning) - is always the Client Operator? or if the Resource Owner grants access in order to make the Requesting Party be another Data Controller, what does that make the Client Operator in this case?
5. **Requesting Party** - 2nd-tier Data Processor (has privacy framework meaning) - is always the Requesting Party? - or maybe the Resource Owner wants to make the Requesting Party be another Data Controller?
   a. **Access Beneficiary**??? - a party that the Requesting Party is acting on behalf of, either themselves or another party - a better name for this?

We definitely don't like the word "ownership" of data or records. Is there licensing of data, or a record, or what? Since the access grant could be about disclosure of data, or deletion of data, or execution of an algorithm, or whatever, the licensing similarly could be sophisticated. It could specify that the recipient gets a copy of the data, etc.

We want to have a concordance between all of these roles and the UMA flow.

Do we use the role terms **Licensor** and **Licensee** at different points in our clause text vs. the terms above? Or can we punt on that for now? E.g., maybe the Resource Owner becomes (dons a role of) a Licensor at a certain point in a contract, and a Requesting Party becomes a Licensee. Two of the parties listed above would be parties to the license/contract, and others are subsidiary. But we still need to complete our mapping work to understand the precise relationships of them all, wrt UMA.

The Resource Owner is the one that has to hire/contract with the ASO, because it can't be someone who doesn't have legal capacity etc. Or even if the Data Subject contracted with some agent, say, an attorney, to take care of financial business or whatever for them, then that Resource Owner really does function as their delegated agent and contracts with the ASO on their behalf. Of course, just as in Google Docs, the agent might need to "transfer ownership" of resources back to the data subject at some point – and the IRM group's design guidelines include just such actions! Delegable, Transferable, etc.

This justifies our "indenting" the Data Subject role; it's "offline" wrt UMA, and our model.

2017-06-16

- Reviewing pieces of deliverable #3
    - All the roles in the use cases
    - Names and definitions for them
    - Different ways to connect the legal and technical artifacts and why
    - The different "solar systems" in the jurisdictional universe

Attending: Eve, Andrew, Kathleen, Thomas, John W, Sal, Tim

Assumption: The resource protected by UMA is personal information. Is it always? We'd had conversations in the past about other use cases, such as that the resource is not PI but rather "confidential information", or "IP", or something else that's not "PI". This subgroup's mission makes clear that it's about "protecting privacy rights", so it suggests that other use cases, such as protecting these other kinds of information (which is useful for both natural and legal persons as resource owners), are out of scope. We should state this somewhere in the final deliverable – not that the framework couldn't reasonably (or even easily?) be extended for these purposes, but we haven't chosen to focus on them at this time.

Data Subject shortens to DS (vs **Resource Subject** and RS), and then it wouldn't clash with **Resource Server**/RS. **Data Subject** is the canonical name in GDPR, the ISO privacy framework, and pretty much all "privacy talk" around the "trinity" of data subjects, data controllers, and data processors. Let's make it so!

If you have a **Client Operator** (say, Google running a **Client**, Google Calendar) and a **Requesting Party Agent** (this is our current term under discussion) – say, Dr. Bob using the Google Calendar **Client** – who is using it under the control of the real **Requesting Party**, a hospital institution – is it abundantly clear that the CO is distinct from the RqPA? Here is how the RSO and the CO each have a legal relationship with the ASO distinct from other/higher-order relationships:

- The RS gets OAuth client credentials from the AS, and there is an opportunity (currently nearly always taken) for the ASO to impose an agreement on the RSO before issuing these credentials.
- Similarly, the C gets OAuth client credentials from the AS, and there is an opportunity (currently nearly always taken) for the ASO to impose an agreement on the CO before issuing these credentials.

When we wrote the old Binding Obligations clauses, we put clauses having to do with the client credentials stage out of scope. However, UMA2 requires the client to take certain actions at registration time (meaning, before the requesting party – NOTE: or, requesting party agent) is wielding that client software. This means effectively that the client's interactions at that registration juncture are on the "legal stage" for us, licensing-wise – so we seem to need a new table for client credentials for each, even if we may not need model clauses for one or both. The requesting party agent note is about the fact that Alice the patient might give access to "a doctor's practice" overall to get access to her calendar for scheduling her pre-op appointment. The receptionist ultimately gains access under this policy. The receptionist isn't "the practice", but an agent of the practice doing work for hire. The calendar is being used by the receptionist (**Requesting Party Agent**) but is "operated" by Google the company (the **Client Operator**). The idea is that the CO would be bound by any license formed through the client credentials issuance timeframe, the RqPA would be bound by his/her employment agreement with the "true" RqP, and the "true" RqP would be bound by the end-to-end RPT licensing through the main UMA Legal framework.

The Word doc uses "Consent License". Why not "Use License" or just "License"? In order to market the framework as being able to solve a large class of consent challenges, it seems we need not put an adjective on "License" – just talk about our solutions in other ways and parts of our framework documentation. Is it worth talking about "Data Protection Licenses"? Or maybe this is a "licensing framework", and we ultimately produce a "licensing toolkit", and hey, maybe a "licensing API" that can be called! The agreements others produce with it could be whatever kinds of agreements they need to produce! Given that the agreements others produce might be dynamic (while versionable etc.) and dynamically negotiable – while we reduce the friction to making them protective of privacy rights – we don't necessarily want to apply old-fashioned labels too soon.

Resource Owner vs **Grantor** as a name is still an open question. HL7 uses Grantor in a way that seems to only apply when the person is the offeror (?) of license terms. But in UMA, this person could be either the offeror or accepter of license terms, based the UX. Or is that true? This is an open question. Eve is starting to suspect that, regardless of flow, the RO/G is always the offeror.

If Eve's assumption is correct, the next open question Eve and Tim had discussed is: Despite the "chain of agents/legal representatives" that may exist, does the offer of license terms always rest at the Data Subject somehow? E.g.:

- **Data Subject** little Johnny (2yo)
    - ...has an Agent/Legal Representative (she's probably the latter because Johnny's a minor) mom Alice, who sets some policies on Johnny's behalf but also
        - ...chose an Agent, Authorization Server Operator ShareHub, which sets up some default policies and security protections
    - Alice shared Johnny's EHR in selective fashion
    - ...with RqP research organization NYP through CO *autonomous web client* Sync4Science Scraper
        - ...which gives downstream access to employee researcher RqPA Charlie

Alternatively:

- Data Subject little Johnny (2yo)
- ...has an Agent/Legal Representative mom Alice, who sets some policies on Johnny's behalf but also
    - ...chose an Agent, Authorization Server Operator ShareHub, which sets up some default policies and security protections
- Alice shared Johnny's EHR in selective fashion
- ...with researcher RqPA Charlie using *web or mobile C app* Research Viewer
    - who gives access to the broader research organization because of his employment relationship

Is the transfer of "access grants" or of a "right"? And do we stick with "Grantor", do we move back to "Resource Owner", or change entirely to something like "Data Subject Proxy" or "Data Subject Agent"? This person may have relationships with a lot of ASOs, and in some cases it's on their own behalf, and in others it's on behalf of others. There is no UMA artifact that documents on the wire (maybe on the wire in the medical world, in some circumstances?) the relationship between these two roles. See our matrix from the 2016-03–18 telecon for all the relationships and technical artifacts. The way the NZ government solved this was to invent a "headless" account for the offline person, and then have the institution manage the account on the person's behalf.

Eve's suggestion for a followup from Tim: Enumerate and recommend answers for all the outstanding questions above.

2017-06-09

- Reviewing pieces of deliverable #3

Attending: Eve, Colin, Tim, John W, Mark L, Kathleen

The new matrix maps UMA consent licenses to key concepts in various regulatory regimes. How can we vet the correspondences? Can we see UMA as another chapter in the millennia-old commercial legal system?

"Open" licenses are freely transferrable. "Limited" licenses limit the use to one specific user. Tim threw in duration as well, under "Non-Transferable". John points out that stopping the usage at one stop could be problematic because then a data controller couldn't ever go on and share with a data processor.

Scenario: RO Alice chooses to share certain files (say, scanned receipts) from digital file system service Dropbox (an RS) to accountant (RqP) Bob, where Bob is using client app TurboTax. (Alice also happens to use TurboTax sometimes.) Alice's AS is ShareHub.

Eve's question is: Is Bob a Data Processor? Or is Bob another Data Controller? Could there be different licenses for either circumstance?

Regarding our artifacts:

- The PAT (where Alice's AS and RS, ShareHub and Dropbox in this case, get associated together) is an OAuth access token, so it's an *opt-in flow* where Alice (typically) has no choice about the terms. The reason this may matter is that we might have limited ability to affect *OAuth's* parameters of a "legal framework", vs. UMA's.
- On the other hand, the RPT (where Alice's policy conditions dictate the specific permissions and their durations) has a more discretionary nature. She has choice and control, to the extent of the AS's policy condition capabilities. (Those capabilities are in the "competitive space". See UMA Grant Sec 3.3.4, second Note text. This allows AS's to compete on policy condition handling.)
- The PCT probably doesn't have any of the same categories that the PAT and RPT have. It's a simple token, actually a OAuth-like, which would have different rows.

Is there *always*, *sometimes*, or *never* a transfer of accountability where the RqP becomes another Data Controller? The use cases we collected earlier, to her mind, seem to include both kinds. Mark brings up the case of consenting to sharing data for marketing purposes in this context. Kathleen thinks consumers won't get subtleties (and Eve agrees). Getting alerts would be a good pattern in case of concerns.

So, apparently, our analysis seems to hold! We'd like to vet all this against the columns of the matrix. ("Consumer Data" refers to the body of commercial law.) Including the parental consent pattern would be good as well; it doesn't have to be specific to any one law/regulation.

- Our presumption is that UMA is about the Resource Server Operator being a Data Controller serving the Data/Resource Subject (under some complex regime of the Data/Resource Subject's proxy, the Resource Owner, and their agent, the Authorization Server Operator).
- We think we have to cover both "sharing with/delegating access to" a Data Processor (someone with limited accountability/responsibility/liability because they're getting access under your control) and another Data Controller (someone with full accountability/responsibility/liability because they're getting access on their own recognizance). We assume that AS and RS competition on interfaces for policy, resource types, etc. will need to ease paths for resource owners who interact with them.

Kathleen asks: How to distinguish a license from a contract? Traditionally, a data subject isn't really given the opportunity to control what's done with their data. Eve: A license is a kind of contract/agreement, right? Mark notes that the CIS group today did a lot of great work around the FIPPS. They stress "... with the consent of the data subject". Discussion ensued about the role of the original principles, further operationalization in the DPD and then the GDPR, and the concurrent rise of the OAuth, OIDC, and UMA stack and Consent Receipts. We also discussed the Brave Browser and their recent huge round of funding.

Kathleen provides a link to a presentation from HIMSS about VA and UMA.

AI: Eve: Send out links to the current state of the (old) Binding Obligations and (newer) CommonAccord model clause text, which was intended to fill in PAT and RPT license text (and on and on), respectively.

AI: Eve and Tim: Meet briefly to brainstorm what rows the PCT would have, and external experts to reach out to.

AI: John, Mark, Kathleen: Review the X's in the cells of the matrix.

2017-05-26

- Reviewing deliverable #2

Attending: Eve, JohnW, Kathleen, Tim, Andrew, Scott D, Mark L

Eve shared the experience of presenting to the Cloud, Big Data, and AI legal conference on May 23. She got to try out the elevator pitch for our nascent legal framework and the rest of the work. She'll point to the PDF of her slides when available.

The first key element of this deliverable, as discussed previously, is the conceptual orange/blue chart. Communications are required in an electronic protocol context. The Communication/Autonomy nexus includes Consent(Autonomy/Law), Sharing(Autonomy/Commerce), and Protocols(Autonomy/Communication).

Scott notes: "Consciousness" as part of this might sound "frilly", but as part of law it's actually not – it's a "meeting of the minds", so it's fundamental to contract formation. There are famous cases every law student learns regarding this concept. John comments, channelling Doc regarding contracts of adhesion: How does this play in? The problem of the human mind and questions of human authority are fundamental. We ultimately have to map UMA to law. Eve: Going by the matrix, this looks like Consent (Law/Autonomy), Access Authorization (Law/Reciprocity), License (Law/Objectivity).

Reciprocity supports enforceability of transactions.

How do we get to a grant of access rights? UMA's permission tokens are an anchor. The word "delegation" from a legal perspective doesn't make sense, quite. Kathleen observes that DRM is a technology that is literally used to effect access licensing already.

Is it possible to get to the point of breaking the law by accessing some digital resource? Who controls the access relationship? Any requirement for consent is a basis for the authority we are looking to ensconce. So the various regulations that strengthen consent requirements, such as PIPEDA, GDPR, PSD2, and others (such as "patient right of access", which is an "economic clout-oriented" strength, as Kathleen points out) strengthen our ability to make these connections. Eve thinks FIPPS is too "soft power", in that it needs to be operationalized through other regulatory structures.

(Eve reminds us all to focus on the "hard power" we have, which is to develop the framework and toolkits to influence those actually deploying services. Our "soft power" to influence policymakers to make law and regulation is to be done only through the framework level as interesting written material, and through other means. That's why we don't have a Resource Regulator role formally.)

Scott remarks that "consent" is sort of a legacy system that isn't fit for purpose, but is well understood. He also notes that the word "control" is a difficult one; a "co-management regime" over digital resources is a more realistic way of seeing the challenge.

Kathleen points out that consumers have the power to fill their profiles with false information; she calls it discombobulation. Eve points to The Economics of Privacy, which provides evidence that savvy consumers can manipulate businesses/organizations, and that greater data sharing can benefit (e.g.) patients.

Eve agrees about stepping back *two* steps from "data ownership": from "own", and again from "control", to "management". *Also* from "data" to "(digital) resource" (because sometimes the data is very transient, e.g. provided through a streaming API). Scott adds that a person's relationship to a thing should really be a relationship to another person vis a vis the thing.

Eve wonders if we can publish a Talmudic-style commentary version of deliverable #2 (or turn all the deliverables into eventual WG deliverables, or whatever) for benefit of the wider audience we're going for. Scott notes that the UCC Commentary provides a potential model; it's citable. Mark notes that trade associations are trying to figure out how to develop codes of conduct that are GDPR-conforming.

Eve points to A Typology of Privacy, and specifically its typology of the objects of the rights to privacy, as great additional commentary that could extend our analysis of Communication; it puts "mediated communication" under the "semi-privacy zone" (which came from Westin's work).

John points further to the NIST Privacy Engineering work; see page 17. Similar to the security concept of CIA, it proposes "predictability, manageability, and disassociability".

Regarding the opportunities to combine forces and possibly have a KI-wide Legal WG that at least encompasses the scope of the current UMA Legal and CIS legal-related work: There's consent receipt lifecycles, User Submitted Terms as ready-made licensing terms, a CR version of the orange/blue matrix?, and possibly more. Eve noted that the current UMA Legal mission is very much like a charter.

The French court had said it was a derogation of human rights to give too-broad permission. (And GDPR has now ensconced a requirement to give specific purpose of use!) So Scott had suggested at the legal conference that maybe it could be possible to delegate to a fiduciary layer some broad ability to manage permissions, and then that layer could give more reliability around this. This was the point at which Eve leaned over and whispered, "Like an authorization server?" 🙂 (Or the authorization server operator, to be precise.) The theory Eve has had is that a Resource Owner (Grantor as was) and Authorization Server Operator should be able to negotiate the former delegating to the latter the ability not just to execute to the former's protection policies, but to set policies on their behalf, e.g. setting default policies. This is a real use case that has come up in HEART, and we have even defined one of the profiles to account for this. Tim says this is where the word *agent* comes in, so it should be covered.

Next steps: Everyone to review the document as sent in email by next week for deliverable #3 purposes.

AI: Eve: Review deliverable #2 for copy-editing/Kantara approval purposes ASAP.

2017-05-12

- Reviewing draft deliverable #2 – which is now a proposal paper!

Attending: Eve, Tim, John (the koala), Mark

Tim showed us his nearly-ready latest deliverable. It's sort of deliverable #2, but it also looks a bit like a full-blown #3 because it's a proposal for a legal framework.

Elevator pitch: The User-Managed Access (UMA) technical protocol applies protection policies to permission tokens. The UMA legal framework maps those permission tokens to licenses as legal devices. This licensing mechanism is valuable to individuals, organizations, legal professionals, and privacy professionals because it allows Alice to license Bob to use her digital resources on her terms.

What do we do about the technical terms that we get from OAuth, "authorization grant", "authorization", "access" (which is colloquial, Eve believes), etc., and the legal (usually, because of regulations) term "consent"? Then there's "permission", "approval", etc. "Grant" comes from OAuth but is also, a verb, pretty neutral. The usual formulation in most regulations is "notice and consent" as a workflow.

Mark observes that "consent to authorize (access)" could be a viable way of seeing how regulators see it. John calls UMA "an affirmative action to allow access". Eve cautions that UMA does not *prescribe* a user experience, and it *might not* involve an affirmative action. Would establishment of a default policy, by an AS and not by an RO, count? That sounds like an opt-out flow.

Eve runs through the Open Banking interpretation of consent and authorization. (See this blog post.) PSD2 has terminology like this, mapped to OAuth /OIDC overlays:

- ASPSP is, roughly, the bank (it acts as the AS/RS – in OAuth/OIDC they're in the same domain)
- AISP (Account Information Service Provider) is like Mint
- PISP (Payment Initiation Service Provider) is like Amazon, selling you something
- TPP (Third Party Provider) is a client app fronting one of the last two guys
- PSU (Payment Services User) is the resource owner

What they require for consent is that the PSU must "give consent" to the TPP (client), not the ASPSP (authorization server). This happens in a fairly complex dance that happens over two stages, concluding in an authorization code flow. So this sounds an awful lot like Mark's formulation. 🙂

BTW, the OAuth spec slipped up and accidentally used the word "consent" once.

It's a goal to make sure regulators, policymakers, and deployers can get comfortable with the notion that UMA deployments can meet "consent" requirements (as long as user experience requirements are likewise met). UX is outside the scope of the pure technical layer, but the mappings we do here can surely remark on UX where we see fit.

Tim plans to finish working on the paper over the weekend. Eve will upload it as soon as she can. Mark is doing mapping exercises in CIS WG; we should liaise on that. Also, note that the BSC DG has put in a recommendation for Kantara to consider an org-wide Legal WG, since this work is going on all over the place!

2017-05-05

- Reviewing draft deliverable #2

Attending: Eve, JimH, Tim, Adrian, JohnW

Peer-to-peer models and/vs. trust framework models, and potential impact on our eventual toolkits:

We had concluded last week that more "onesie-twosie" (peer-to-peer) contracts could be appropriate in many more circumstances than "trust frameworks". It takes a lot to push ecosystems over into a fairly expensive and static agreement-making process (as seen in identity trust frameworks vs. regular contracts), and as the BSC analysis showed, that often doesn't favor individuals anyway. A finer-grained process *that reduces friction* and achieves the goals of our subgroup would be helpful.

Jim points out that there can be a leveling aspect of one-to-many. In both 1:1 and 1:*n* cases, the key is *creating efficiencies* in the direction you want. This may speak to the kinds of tools that we'll want to develop next based on the legal framework that we develop first. John suggests that maybe we want to create "Binding Transactional Rules" as a kind of tool, analogously to Binding Corporate Rules that are intended to be static and 1:1 between enterprises. Nice! If they can be win-win, both friendly to business and friendly to the business's users, and also more dynamic, they could be successful.

Jim notes that the IACCM tries to move towards supply chain and infrastructure *relationships.* A master services agreement is a key framework artifact, with layers of parameters on top of that.

If we can add individuals to such relationships, and even allow for "one-night stand" relationships 🙂 (that is, say, a single purchase or payment or access or whatever), then a way of standardizing terms that the individual insists on could feed into a contracting workflow such as this CommonAccord example. This is one UMA Legal toolkit use case: model clauses, at least for purposes of granting access to digital resources. Would reducing friction (cost) to using such tools be enough, in the absence of business-model incentives, to encourage a resource server to give a resource owner meaningful control?

Other use cases would be for others to develop their own legal toolkits, as HL7 is doing – described in the BSC report. (An UMA "technical" use case could be delegating access to such contracts.) A hash of a signed contract could be recorded on a blockchain.

The licensing model, completing deliverable #2, and "the three tokens":

How can licenses be converted into tokens? We have some urgency around figuring this out, given the regulatory situation – not just GDPR but PSD2 as well, and even the HL7/FHIR and perhaps DIACC/TFEC/PCTF work going on.

Would this be like a "reverse EULA"? Adrian's work with Patient Privacy Rights has revolved entirely around the license concept. The patient shouldn't ever have to see the CommonAccord details we've been looking at. If you look at Creative Commons, everything is given a name/nickname, which helps. And the User Submitted Terms work tries to do convenient bundling as well.

John describes "the tyranny of the default". There is great power in default settings. We can reduce friction/cost in "doing the right thing". Jim provides another MSA example where some default choices have been made.

We do have a call next week. Let's plan on reviewing a final (barring final amendments being suggestions) deliverable #2 on May 12. That call will have a hard stop so that the full WG can start meeting right afterwards.

2017-04-28

- Reviewing draft deliverable #2
- "Resource Regulator" role

Attending: Eve, John, Adrian, Kathleen, Tim, Mark, Bjorn

Review of Move Fast and Break Things, recommended by John. (We were talking about turpentine and the importance of something stinging in order to tell that it's working – also "if it hurt, do it more often" from Agile.)

Resource Regulator role:

How would we include Resource Regulator in any charts, if we did? A regulator constrains, by law, what some of the other parties do, in a context of enabling others. In other words, it sets rights and responsibilities. And we can't change these. These are jurisdictional. So a regulator is a sort of "god". It's out of band of the UMA protocol and Eve suggests it's also "out of band" of UMA legal because we can't affect regulatory/public law on any reasonable time scale.

Here's what we're doing in this subgroup:

1. Building a legal framework...
2. So that we can next build "toolkits" that are friction-reducing building blocks for various contractual mechanisms ("private law")...
3. So that third parties can deploy UMA-enabled service systems in a manner consistent with protecting privacy rights using those contractual mechanisms that adhere to the laws and regulations ("public law") of their jurisdictions

Kathleen points out that in the case of HIEs, e.g. in Michigan, there's a kind of blending of private and public law – there's a contract made among HIE participants that includes 42 CFR Part 2 by value. But that was their choice; by-reference would have probably been better. (Hey, CommonAccord would have helped. 🙂)

Eve's proposal: We can influence interpretation of regulations (especially now when a lot of GDPR etc. guidance is being written), but let's keep "regulator" lowercase since we can't influence written regulations and laws directly. **Consensus** not to uppercase the Resource Regulator party role.

Roles and splitting:

We have roles that mostly stick to the existing UMA technical roles, with the exception of **splitting** Resource Subject and Resource Owner for now. Once we have some model contracts, we can add nuance in splitting roles. Our previous "party" terms of Authorization Server Operator, Resource Server Operator, and Client Operator weren't about splitting but just about acknowledging the different between "software entity" terms and "party to a contract" terms.

Trust frameworks and access federation:

The only ones we are familiar with that exist (that could potentially accommodate UMA-style – party-to-party – user-centric authorization and "consent directives", vs. just opt-in consent) already are trust frameworks for identity federations, except for the emerging Pan-Canadian Trust Framework and the very new health trust framework efforts that Kathleen has mentioned in the BSC context. The latter work puts the data under the patient's control under the "patient's right of access". There are liability shifts which would result in the motivation becoming stronger for appealing to individuals to ask for data for secondary-use purposes.

Interesting situation: For a genome, there's more than just one data subject for the same data set!

(Should we be talking about reducing friction in "federating access" in UMA-enabled contracts and not yet talk about automation of true "access federations" a la Shibboleth-like automated onboarding of IdPs/RPs?)

Licensing as our legal model:

We need to use some well-understood legal model in our "private law" tools. Effectively there has to be some sort of cascade of rights that the RO gets. As Kathleen points out, the RS will only expose licensing rights that meet their business model. We hope that reducing friction to "doing the right thing".

Tokens:

There are three tokens that come into play:

- Requesting party token (RPT) - required
  - The RqP is able to revoke this through the OAuth token revocation mechanism.
- Protection API access token - required
  - The RO is able to revoke this through the OAuth token revocation mechanism.
- Persisted claims token (PCT) - optional

- The RqP is able to revoke this through the OAuth token revocation mechanism.

2017-04-21

- Reviewing draft deliverable #2

Attending: Eve, Tim, John, Mark

Tim's insight around identifying the "harms" to the parties in the #2 exercise helped guide the development of the draft deliverables we're looking at today. John opines that this view elides the "rights" basis for privacy breaches because it's property-based. Well, this is the question. What can we effectively achieve with our clauses and other tools? If agreements/contracts are the basis for what can be achieved between/among a resource owner and other parties, what are all the choices for legal theories? Tim is proposing a licensing basis. (We discussed this back in 2017-04-15 and seemed to reject this, but what are other alternatives?) There is a governance function and also an economic function.

Looking at Sec 2.1 of the EDPS opinion on digital content, John points to some commentary on the VRM list where someone was troubled by the "market for personal data". The point they were making was that someone could agree to selling organs (or their body into slavery or whatever), but this shouldn't perhaps be possible with selling data. We in UMA take a different, more empowered/powerful, position.

Tim's Chart 1 is more of a windup to chart 2, and he will supply more explanatory text for it. The "Communicative Behavior" column means how the requirements for Value, Meaning, and Information are conveyed/communicated, e.g., trust frameworks, regulations, configuration documents, API documentation, etc.

Both are about the relationships formed, and are explicitly not about "data ownership". Chart 2 is the "money chart". (Eve screenshared them, and Tim will be revising these and making them available to all before next week's meeting.)

So can we state the following?

- The data subject has rights over the information about them.
    - True as part of the Universal Declaration of Human Rights.
    - Different jurisdictions ensconce this right to different degrees in law/regulation or not.
    - True of information even prior to its being digitized.
- The data controller and the data processor have property rights related to records containing a data subject's information.
    - The records could be in digital form or not.
- The formal "interface" (communicative behavior) defined between data controllers, data processors, and data subjects is regulations.
- UMA has the potential to enable data subjects ("resource subjects") and their proxies (resource owners), or even data subjects on their own, to consent to data ("resource") access by third parties ("requesting parties") in such a way that the third party is a data processor.
    - We believe the regulations are currently blind to:
        - The proxying opportunity in UMA
        - The potential ability for UMA to distinguish between granting access to someone who fills the role of a "data processor" vs. "another data controller"
    - UMA only has soft technical constraints (the "Adrian clause") around jurisdictional nonfunctional requirements for things like data localization.
        - The potential extension for "cascading authorization servers" would provide a potential hard technical solution.
        - We have the potential for providing legal toolkits that give legal solutions that may suffice.

Do we need a Resource Regulator role?

If you're interested, there is a SAMHSA Consent2Share webinar on April 25 at 3:30pm ET. Registration link is here.

2017-03-24

- Reviewing use cases doc

Attending: Eve, Andrew, Tim, Adrian, Paul

NOTE: No Legal call next week – and possibly the week after!

AI: Eve: Look for substitute Legal call as necessary.

Tim reports that he's well on his way with deliverable #2.

Let's take a "breadth-first" approach to the use cases in deliverable #1 today. Tim asks: What's the general nature of the authorization services being developed today? Are there any? Are they primarily C2G or C2B? Yes, and they are split, with also some B2B. Eve ranks the top three-to-four use cases as healthcare (and healthcare-plus-IoT, with IoT generally being quite broad and horizontal), government/citizen interactions, and financial. Adrian notes that regulation is a powerful motivator for using UMA, and thus the consumer component seems to be critical to adoption.

Use case C: And here the FCC decision from October just got reversed again! Reasons to think this use case for UMA is valid: UMA would give a standardized interface for ISPs to offer third-party partners for efficient integration (more money on both sides), and wherever regs allow either "legitimate business interests" or "consent" as valid bases to collect and share data it's better to ask for consent because of user trust risks and creeping regulation scope (IT cost and user trust risks of having to ask for consent later). This could be Alice-to-Alice sharing (she shares with some app that she will use) or Alice-to-other sharing (e.g. for marketing purposes, maybe even "targeted" for some goodie offered to her, but maybe not).

Worked on use case D.

Worked on use case E. We decided to make it just the parent/guardian use case.

AI: Eve: Try to add a payment use case as a variation on D.

2017-03-17

- Reviewing use cases doc

Attending: Eve, Tim, Colin L, Adrian, John W, Ann

Sal provided info on an event being held Mar 23 on Information Law. We also discussed Jim H's Wise Contracts paper and Jane Winn's comments in another context (BSC). This is relevant to Tim's intended matrix regarding allocation of economic value.

Adrian sent a link to the Google DeepMind problem with trust, through extreme lack of transparency and communication with overseers. Eve's "golden rule" is to tell business owners that they should treat personal information as a joint asset. John W tells them the business owns the record while the individual controls the information. Note that OAuth's use of "resource owner", and thus UMA's effectively means control of access (to some extent/scope of access). Adrian "is alone in using" the definition of ownership where you can delete it. We've stayed away from "owning" terms so far for all these reasons, except at the technical level.

Tim is listening for device connectivity, issues of content, and issues of context. For example, how to prove the connection of a device to a responsible mind? Eve points to a particular OAuth grant flow called OAuth Device Flow that helps to bind a device to a person and their account, which could help in an UMA context. Should we include a connected car use case? Here's an example we could potentially use.

Regarding the distinction about the legal bases for collecting and using data (and also for presenting interfaces and granting access beyond just "collecting and using data", keeping in mind that UMA can protect any API – think "resource controller" vs. just "data controller"), is this exactly the bright line that lets us say that this lets the resource server be a "resource owner" for those resources that it doesn't give Alice the rights to control access to? This may be tautological because UMA has a notion of enterprise ROs anyway.

Bridging terminology thoughts:

- From resource server to resource server operator to resource controller to data controller????

2017-03-13

- Reviewing use cases doc

Attending: Eve, Adrian, Tim R, Paul L, John W, Mark L

The "Use Cases Deliverable 1b" document (in GDoc form) is now basically trapped in amber, so to speak, and the subgroup might tweak and change for our own purposes, e.g. starting a new version of our primer materials with this content, or editing this doc, but Tim will probably create the other two deliverables owed separately from this one.

AI: Eve: Convert the "salient factors" to hyperlinks in the doc.

The doc starts with legal considerations and use-case desiderata, moves to proposed use cases (which are "hybrid" in that they are as real-life as possible and mix various scenarios), and finally moves to UMA technical considerations that need to be solved. Today, let's analyze the use cases with the thought in mind to figure out whether there's a finite or infinite number of them to capture.

Should we "go there" when it comes to things like industrial IoT (IIoT or Industry 4.0)? We probably have to because the topic of national infrastructure and mass transportation with smart sensors has been coming up, and individuals interacting with these, with privacy implications. So we should probably include use cases for them.

We got through use case A and part of B. Tim is keen to keep use cases C and F, which might not at first blush seem directly applicable, because of the market need.

AI: All: Review the use cases doc (asking Eve for commenting privileges as necessary) in preparation for Friday's call.

2017-03-03

- Reviewing draft Legal deliverable #1

Attending: Eve, Tim, John W, Colin, Mark L

Logistics:

Since Fri Mar 10 Eve can't meet, we will substitute Mon Mar 13 at 10am PT/1pm ET.

AI: Eve: Schedule the substitute meeting on the calendar.

Legal deliverable:

Tim has deep experience in identity and authentication law (Virginia and unified US, ++). We are fortunate that he is a passionate legal UMAnitarian!

The first section in the doc is about goals. User "management" is about rights to *control* access. An ABA publication came out just today about IoT, addressing topics of data "ownership". Mark L talks about the distinction between Data Protection and Data Control, and UMA is really about reaching the latter. UMA goes beyond most of today's regulations, and that's what's groundbreaking. Tim believes "this could be the most significant legal tool in a thousand years".

Access control can have governance and economic functions. Data flow can unleash positive value. Business value and individual value can both be served in tandem by enabling selective sharing.

"Diachronic" consent/access control is about allowing ongoing changes and additions of information. Value can increase over time as information is shared. So this concept is directly related to business and individual value.

AI: Eve: Convert the use case document to GDoc form so we can comment/adjust live.

The *Lex Informatica* considerations are about ensuring that our global networks and global flows of data are accounted for. There need to be rules embedded in software and devices.

For the "F. Citizen-Facing Government Services" use case, what if we develop an additional use case that matches the OIX Pensions Dashboard use case, which is similar and another live topic of conversation?

Parking lot of topics for Tim and the group:

- Distinguishing between cases where there's a *hard statutory reason* to disclose data/give access (e.g. GDPR Article 6.1, HIPAA, etc.), and not giving the individual an opportunity to consent and say no (where access is given anyway), and cases where there's an organizational decision to reserve to itself data sharing rights with an opt-in that's not "data control"-friendly
- Adding "UMA salient factors" more specifically in the section after the use cases, including possibly proactive "Share" vs. reactive "Access Approval" (Opt-In with Choice) flow options

AI: Eve: Share RSA talk link, consent layers diagram, the article with the articulation of "no data ownership", and snippets from the article on the Internet of Medical Things.

AI: Eve: Reach out to Mike Pegman/David Rennie about possibly developing a companion use case for use case F.

AI: Eve: Reach out to the people who agreed to review the primer a long time ago, and ask them to look at the deliverable as it grows and see what they think (maybe enhanced with our "UMA technical definitions".

2016-12-16

- Use case/mapping exercise
- Toolkit discussion: roadmap?
- Legal as a WG vs. a subgroup

Attending: Eve, John W, Ann

Could a PIA be a target for a "toolkit"? John does a lot of these. There are lots of standard forms. ISO is working on one. A PIA is a tool for evaluating the policies and practices of an entity wrt a given standard (law, directive, etc.). So to the extent that you can log that a particular technical interop standard such as UMA or OTTO or whatever has been used, it's likely that this is necessary but insufficient. Under accountability, you should appoint a DPO, you should have contractual elements so that transferring data to a processor has valid and correct safeguards, etc. Eve is imagining that if we have produced some model text and it's versioned, jurisdiction-specific, etc., then an entity that actually uses it as a "toolkit" is in a good position to also use a subsidiary "PIA toolkit" in efficient fashion to get a leg up on compliance. Our toolkit work is precisely intended to be the "BLT stack" that fills in enough real-world deployment pieces above the technical standard to be interesting.

Make this a separate WG vs. a subgroup? It's effectively a parallel group already, without the logistical overhead. It does want to produce "real specs". Our intent is to boil a smaller pond if anything rather than open up to new charter implications, while still being open to the appropriate extensibility points that would allow for things like consent receipts (and maybe other similar proofs of consent), UST (and maybe other similar), etc. to be referred to. The default is to keep everything the same for now.

The axes of use case development actually look more like this:

- Data subject vs. some proxy for them granting access
  - Granting access to a resource that is not individually but jointly "owned" in some fashion (e.g. genome)
- Complexity of party that was granted access (employee of organization, individual acting on own behalf, organization itself...)
- Task-based access grant or control-sharing access grant
  - Does anything about this change if some "non-HTTP-GET" operation is used (i.e., the client and requesting party aren't being "disclosed to" but are actively operating on the server's contents – inserting or deleting)?
- Fiduciary role of the operator of the AS (SaaS, personal, enterprise)
- Jurisdictional location of the grantor/data subject/grantee (with any complexity)
  - Matters to operator of the RS for their data transfer accountability
  - Might matter to operator of the client for their data transfer accountability if "non-HTTP-GET" scenarios matter?
- Sector-specific and other use-case-specific details

The UMA scope mechanism isn't a good match for purpose limitation because the RO would want to differentiate "user-submitted terms" per grantee. But some factored-out legal text might actually be appropriate to add to scope description documents as metadata, at least in the case of "disclosure-oriented" scopes, depending on how our analysis of use cases comes out!

We still have toolkit roadmap work to do. That would be appropriate after more use case/mapping work.

2016-12-02

- Use case collection and impact on concept mapping
- End-of-year planning and news

Attending: Eve, John W, Kathleen, Ann, Paul

Eve spoke with Tim Reiniger this week, and they contemplated his being able to help with our mapping and use case gathering exercise. ISO 29100 (free to download – John will provide the download site), an international standard, has definitions of the actors and roles – so this is particularly useful for our generic purposes. (Note the scenario table. Scenarios g and h involve third parties, so those are about something like "downstream" access/sharing /disclosure.) So our current thinking is that we could make good progress on this exercise in short order. We looked at GDPR's Article 6 to explore its rules for enabling a processor to do its stuff.

We're hoping to find use case archetypes that are ones our audience most wants to see solved, which can drive a limited number of mappings. Otherwise it will be hard to get our arms around the "toolkits" we want to create. We actually want to solve an unlimited number of use cases, only we want to do it by cleverly developing tools that are responsive and parameterizable to those use cases, using limited resources.

Thinking of which use cases we want to collect: Having some health and some IoT and some public-sector is good, but we want to spread them around. See these BSC DG notes for three. The NZ POC of UMA has a ton. We should invite all of our privacy eagles and legal eagles to contribute their use cases from personal experience.

AI: Eve: Talk to Colin to get the logistics into the works.

After our meeting on Dec 16, we officially have no Legal meetings scheduled for the end of the year! Eve suggests publicizing ad hoc "editorial meetings" to the subgroup, and everyone joining as they can. That way, we can try and live up to our very new mission statement about 2016 deliverables. 🙂

2016-11-18

* Eve to give update on Karsten Kinast convo

Attending: Eve, John W, Kathleen

Eve had a chance to deep-dive on some of our current topics with Karsten. In the rights-based EU conception, law/regulatory interpretation/contract can capture the sheer reality of little Johnny's natural rights to privacy over data about him; he has an aura that can't be detached. His relationship with his mother Alice is internal, and UMA's relationship is with him, not Alice. This is "the view of the world from the human rights/legal end". On the other hand, "the view of the world from the technical end", where the technical end in our case is UMA, would connect UMA to Alice because she is the one actually wielding the technology as a resource owner. So the bridge we're building connects those two worlds at the point of identifying the resource owner, data subject, and grantor roles. Does that work? For this reason, it's perhaps best to use the literal GDPR (and other data protection law) wording, vs. being cute and using "resource subject".

This is making us want to revise the charter to a) ensure that our work helps UMA be used for good and not evil generally, and also b) ensure it's used *on behalf of* data subjects, not just for resource owners in that role. How about replacing ", particularly where the resource owner is an individual" with "in a manner consistent with protecting privacy rights" (and get rid of the word "themselves")?

They also talked about sharing access with the intent of making the grantee a(nother) data controller vs. making them a data processor. Making someone a DP (say, because you're asking them to do something "on your behalf" for consideration – e.g. outsourcing payroll processing) means accountability remains with the grantor (Alice), while making someone another DC means you've just given them access for whatever reason, and they're just another accessor. Control becomes additive. So Eve's working theory became: Would it be possible to develop tools that distinguish between Alice's ability to grant access for the purpose of processing vs. for other purposes, as a broad means of managing the DS's/ risk? Would our time be best spent in this group working on these tools first vs. other tools? Could grantors/resource owners clearly understand such distinctions? This is where the DS would be accountable but the DP would need to take on responsibilities (think "RACI") to balance the risk. Our tools – e.g., data sharing agreements – could include standard text for that.

The last subject they discussed was the types of use cases that would need to be developed. Eve has started a fresh Legal Use Cases page, here. It's just a shell so far. Following are the use case axes identified so far.

* Granting access to DC vs. DP
* Disclosing data vs. granting other kinds of access (e.g., letting a client of a grantee add data back through the API, or let it use an algorithm by paying for access, or accessing a photo...)
* Whether the DS and the grantor are both in the EU (leave this complexity aside for now??)
* Every vertical and horizontal use case will probably need its own special-sauce text! But we should at least start gathering examples and identifying what are hopefully "design patterns"

John is concerned that XACML went down a similar path, with resulting complexity. Let's seek ways to boil the smallest of ponds.

2016-11-04

* Review draft UMA Legal mission statement (and Jim comments), wordsmith, and decide
* Review current use cases and plan for building up new ones

Attending: Eve, Sal, John W, Mary, Adrian, Colin

Shoebox/notification endpoint: This is getting more and more important for the WG side to work on, for "legal" reasons. John W notes that this is where "the blockchain as syslog for the Internet" 🙂 could have a role, if you don't trust centralizing your audit log or can't centralize it for any reason. The legal agreement between the parties in question could contain a prescription for *how* the notifications/log entries must be delivered. Eve is suspicious of "newer" technologies, even including API endpoints, for this; emailing stuff is what's used a lot today for TripIt, Expensify, and so on, and it's even aggregatable and combinable with "recipe" technologies.

Should the WG standardize on a UX for UMA interactions? Or maybe it should be about a "model UX" to give ideas. Eve is usually in favor of different communities being free to standardize or profile their own UX (like maybe HEART), but sample UXes are great.

CASBs and enterprise federated authorization use cases: This seems to keep coming up more and more as a potent use case alongside "Alice" use cases. UMA has information security and cybersecurity implications as well as privacy implications. We've talked about this in the context of the primer, and so far we've mentioned it in a footnote, just to say that we'd deal with it in a separate document. Adrian recently mentioned UMA having a role in in rate limiting and such. And Eve recently mentioned UMA being a kind of loosely coupled PDP/PEP system.

Mission statement: John notes that the mission is substantial enough that it almost suggests a proper WG of its own, vs. a subgroup. If we can attract to our table additional expertise of the right sort, does the distinction matter? This WG has the right IPR policy for supporting the kind of work we want to do. And for the existing cadre of people who attend, not having enough time in the day doesn't change no matter how many WGs there are. What do Colin and Andrew Hughes think?

Adrian would like to produce a minimal-subset legal agreement without and then with UMA. We were planning to put this at the end of the primer, once we finish our model text building work. And we might need to presume a particular use case in order to provide the example. But it's definitely the best way to explicate the usage of model text.

Jim's comment was: "In connection with the focus on GDPR and common vocabularies might help bring together a couple of threads to experiment with document "skins" on the GDPR and collections of uses, data types, etc. For instance, to what extent can the GDPR vocabulary of "personal data", "processing", "controller", "processor", "recipient", "third party" and "consent" become a general approach or be extended into one." Agreed that these additional terms are key. We also discussed happiness with ongoing progress on the CommonAccord technology, and our hope that methods of timestamping included-by-reference text are being included.

We'll take on this mission statement as is for now, and keep discussing whether it should turn into a charter statement.

2016-10-28

- Working session on User-Managed Access (UMA) in Contractual and Regulatory Contexts
  - Working through use cases per party in the transaction

Attending: Eve, Jeff S, Kathleen

Parking lot

A FAQ would be: What about the identity story? What about where the identity of the RO would be stored vs where the consents would be stored? (thinking of Eve's "scenario 1", which is the UK Blue Badge one)

FAQ idea

We should add FAQs in the doc itself, to answer questions right as they arise in readers' minds, such as:

- Once data is unshared, doesn't the recipient already have the data? What happens after that?
- This model seems to assume that the enterprise doesn't have any overarching policy itself, and just allows a user to have any sharing policy they wish. Is that correct?
  - This relates to the larger topic of moving the setting on the continuum of org/individual control more towards the individual; see "blockchain identity use case"
- What are specific burdens around PHI vs generic PII?

Let's collect these and be as "pointed" as possible in formulating them. We can decide how and where to answer them as we go.

Funded legal analysis and use case work

Eve reached out to Karsten, who likely doesn't have time for the work himself, but who has kindly made himself available to chat with Eve to suggest next steps. She will take action on this.

Eve suggests that we should decide, by the end of the year, which deliverables to produce in 2017 as "toolkits" (of some sort) for *all* of the frameworks out there. We discussed our rationale for this: It's to get from our very early stage of UMA adoption to exactly one evolutionary stage further. 🙂 Meaning, we recognize that organizations have incentives to gather data, sometimes act badly towards less-empowered parties (individuals in the main), and so on, and we are looking to demonstrate benefits to those organizations – particularly business and legal audiences -- of the use of UMA through educational materials and reduce friction in using UMA through our "toolkits" (which could be model clauses, could be BCR tools, could be consent receipt templates or profiles, etc.).

This is basically a further sharpened mission proposal, if you compare to our 2015 and 2016 versions.

AI: Eve: Propose a sharpened mission statement on the list for review.

Regarding the number of use cases in the world: The hope is that it's more like prepositions (a couple of dozen in English) vs verbs (essentially infinite)! But if there are really any number of them, probably we'll have to identify the most common ones that have lots of examples that hew to a pattern, and leave the "long tail" ones alone. Eve is planning to document these briefly in the wiki and prepare them for our hoped-for legal expert to review, so that we can get to a place where our "toolkits" can supply good tools that map neatly to GDPR interpretations.

2016-10-21

- Working session on User-Managed Access (UMA) in Contractual and Regulatory Contexts
  - Working through use cases per party in the transaction

Attending: Eve, John W, Ann

What is the scope of our work (in this subgroup)? The technical layer may or may not change, while the contractual and regulatory layers above may change surrounding it. Our effort to gather use cases is an attempt to literally scope our work for now. Think of this as a "V1" effort of the "UMA legal subgroup" (which perhaps needs a new name because "legal" is too specific?). The use cases are driven by business and influenced by the regulatory environment, and could be implemented with UMA *and* other technologies – e.g., consent receipts, user-submitted terms, JLINC's stuff, etc.

Eve's theory: We wanted to do draft model text by 1Q2016. 🙂 But the need for some visible progress by 4Q2016 is getting acute. Evidence: The use case of "delegation of access rights" (think "power of attorney", or positive acceptance of access on behalf of the resource owner, or "I am accepting that you gave me rights to drive this connected car", or "this is a payment API and Alice gives Bob the right to spend money out of her account") keeps coming up. UMA the technical solution says nothing about the API semantics, but the contractual and maybe the regulatory layer (think Open Banking API and PSD2) may say a lot about this.

Also, there are some complementary technologies such as Token Exchange, which has an "act_as" semantic, which may have a role at the technical layer. Once upon a time, we had innovated the "Binding Obligations" approach, which enabled a mapping of the non-technical layers to the technical layer where there was a state change so that it's possible to capture that change as part of the chain of evidence (hey, with blockchain/DLT involved??). Sometimes there is no technical state change at the UMA level, but sometimes there is, e.g., a token expires or gets issued or something. And we can start to be more and more clever about that.

Keep in mind that "getting access" could mean getting a value of data, or getting access to a stream of data (e.g. Websocket), or getting access so as to insert or delete data. In the world that used to be, only control of "disclosure" mattered. But now, with UMA anyway, API publishers (RS's) are authoritative for what grain of access is possible, and ROs can control what scopes are handed out.

What are the constraints on Bob's behavior once he gets access as a "first-party" RqP?

1. Bob is a completely free agent to use the access as he sees fit.
2. Bob is legally or contractually constrained in how he uses the access.
3. Bob is technically constrained in how he uses the access.

Eve attempts to make the case that there are fine gradations (referencing her as-yet unpublished permission model taxonomy!) among "ways he uses the access", e.g. onward sharing with third-party grantees such as Charlie, and using it for marketing purposes vs. essential business purposes. She owes the group/the world her taxonomy writeup already.

Look at last week's notes for the description of how UMA and a combination of other technologies could address the onward sharing proposition.

The "**delegation of access rights**" use case seems to be such a popular one that maybe we should highlight it in our interim publications.

What are the realistic options for Bob's use or disclosure of Alice's data he's been exposed to (in this specific case, vs. onward sharing)? "Information wants to be free", generally. To the extent that data is volatile, it loses value over time. To the extent that data is nonvolatile and valuable, helping Alice not to overshare will help, and leveraging wider ecosystems where the business model favors Alice's interests. Beyond that, only nontechnical methods of enforcement will need to be exploited to keep the Bobs in line.

2016-10-14

- Working session on User-Managed Access (UMA) in Contractual and Regulatory Contexts
  - Working through use cases per party in the transaction – see last week's notes for the pattern

Attending: Eve, John W, Adrian, Kathleen, Mark, Mary, Sal

Adrian points to this fascinating article introducing the notion of an "information fiduciary". We discussed the merits of this phrase vs. "agent"; in the past, we've had feedback that "agent" is an accurate legal concept, but too academic. Mark concurs. Some give feedback that "fiduciary" may be helpful has a concept. It does have a connotation of totality of trust. This is sitting pretty well with quite a few use cases. What if our model text, if adopted as a whole by an ecosystem, would confer a kind of "safe harbor" (see the reference to this concept in the article) that lets the various service roles declare that they are "fiduciaries" of specific kinds on behalf of the resource owner/Resource Subject. At the least, we could advocate for this to be a standard interpretation among the Pan-Canadian Trust Framework developers, the GDPR regulators, etc.

We're thinking that the budget we have available to use – and must be used by end of year – would best be used in a completion of this analysis.

In fact, do we want to take BCRs and turn them into our text, or go in the other direction, or what?

We have defined **Resource Subject** with our eyes open to the analogy with Data Subject. Does it make sense to have **Resource Controller** instead of Resource Server Operator? This seems relatively close, by analogy to Data Controller. But then again, from a rights-granting perspective (where the DS is the rights-holder on whatever basis and the DC works under constraints set by the DS) and what we've done with the design of UMA, we have perhaps split out the Data Controller role into two with our RS and AS. This is part of the innovation of UMA.

And then what about **Resource Processor**, analogously to Data Processor, for the Client Operator? Is disclosure of data always for some purpose that is "on the DS's behalf", or is it sometimes for the autonomous use of the access recipient? Eve wonders if the "behalf" distinction is just a way of saying that liability, or responsibility, or accountability, is really strong or has a lot of components to it. As long as the regulatory/contractual layer above the technical UMA layer can capture the subtleties of the liability (such as "you can only share this data after a certain time") or even the technical layer can capture some constraints (such as "sharing will end at this time" or "you cannot write data back to the server, only read it from the server"), then the entire requesting side – both client and requesting party (however each ultimately get split out in "party" terms) – could constitute the Data Processor.

What about the challenge of "on-sharing", or re-sharing, or downstream sharing constraints? Here are some thoughts:

- If there is a narrow ecosystem where all users are using the *same* AS, it's very easy to manage this, because the (singular) AS can keep track, and then you (the RS that publishes APIs) could theoretically invent something like a "no re-sharing" scope on your APIs and then let people set this when they share. (Think of Google Apps's similar "Advanced" feature.) Such a scope, if that's the right place to put it, maybe could be a standard "best practice" scope in various APIs.

- In a medium or wide ecosystem, where everyone uses (say) a different AS, John points out that Alice could at least monitor sharing through a ledger mechanism (he can send the diagram to the group), but you'd need something like encryption or DRM technology above UMA to actually control the re-sharing. John mentions JLINC and Sal mentions COALA.

AI: Eve: If feasible, approach Scott D about the possibility of more focused work to complete our use cases and terminology and analysis around DS/DP /DC.

2016-10-07

- Working session on User-Managed Access (UMA) in Contractual and Regulatory Contexts

Attending: Eve, Kathleen, Adrian, Mary

Let's do some use case work today.

Kathleen notes that, working with the US feds, there's a difference between a "contract" and a "memorandum of understanding" (MOU). In an MOU, they're funding a program, so some transaction value is flowing. Many types of contracts do capture some sort of payment. Maybe this has to do with the relationship being formed, and therefore impacts the model clauses most forcefully.

Analogously to the Grantor/Resource Subject split, we anticipate there needs to be a Grantee/*something* split, where there's potentially someone who might request and get access as part of an employment contract on behalf of the real party who's accountable for the access (like the doctor or admin who works for a hospital). Mary brings up "**Relying Party**" as a candidate for the opposite side of the Resource Subject. Eve has some discomfort because it may cause confusion between identity federations and access federations. But what if the analogy with the sharing specifically of personal data is so powerful that this phrase evokes the right images? Or maybe **Receiving Party**? **Resource Recipient**?

RO-side use cases, summarized:

- UC1: RO = Grantor = Resource Subject, where this natural person has legal capacity
- UC2: RO = Grantor, where Resource Subject is a natural person without legal capacity
- UC3: RO = Grantor, where Resource Subject is a natural underage person unable to consent online (yet) and they are also a RqP
  - We need a **UC3a** that maps what happens when the underage person becomes of age
- UC4: RO = Grantor, where ...

What are the RqP-side analogies? Let's try and fill those out next time. We need to build out these lists for every technical role.

AS-side use cases: Adrian's main concern is the AS(O) as an agent of the Resource Subject or Grantor, where there are the strongest possible constraints (technical if possible) on anyone but them seeing their policies and the RS cannot known whether the policies are under the control of either the Resource Subject or Grantor. In this context, he wants to describe the limitations of liability of an RS(O). The issue here is that the service we would call the RS currently has visibility into a sign-off by people in both the resource subject and guardian/proxy roles in some fashion, and since UMA interposes an AS as an "agent", any service becoming an UMA RS would lose this visibility. The legal questions for us are: Is this loss of visibility acceptable? If not, do we have to build a facsimile of the visibility into our model clauses? Are there jurisdictional variations?

AI: Adrian: Add the above as a GitHub issue.

PARKING LOT: One thing we should discuss at some point is the notion of "pushing" content" vs. a typical web pattern of "pulling" content. The usual web pattern we think of is that a client approaches an API unilaterally, which is nominally a "pull". However, there are patterns such as websocket where there is a constantly open channel through which the RS could push content at will. The IoT uses this a bunch, and some other technologies. Kathleen mentioned the AS serving as a content broker – we'd have to discuss what this means more.

2016-09-23

- Working session on User-Managed Access (UMA) in Contractual and Regulatory Contexts
  - Today's meeting is at 10am PT/1pm ET/6pm UK/7pm CET
  - Same bat channel

Attending: Eve, Ann, Scott D, John W, Mary

Eve's slides that show four "BLT use cases" are here. In UC3 (Alice oversees 12-year-old Susie's online usage), there are different kinds of harm that we're trying to protect against. Some are about inbound data sharing, some are outbound data sharing, some have other natures. Also, the diagram doesn't show that Susie is expected to be able to use her UMA client in turn as a downstream UMA RS. In some contexts/jurisdictions, there are rights that Susie would have that need to be protected as well. The "age of majority" subtleties hinge on various ages, and it can be a combination of a legal and a contractual status. We've imagined that it would be a technical status (PAT revocation) that would reify the status change that would allow "timing in" of legal capacity.

Though the primer probably wouldn't go here, and our UC diagrams are static, we could theoretically write UMA profiles that specify "state changes" such as PAT revocations and such that are legal capacity-dependent, perhaps in some complex workflow relationship with the type of the digital resource being protected. This could even enable "smart regulation". The Past Is a Foreign Country – can we overcome our biases about how we've done things previously? We suspect that this challenge is actually operational, rather than heavily theoretical. Once an AS knows the age of a resource owner, and an RS can share well-known semantics about a resource set type with the AS, then policy conditions can connect the dots about jurisdictional constraints on sharing information about, say, pregnancy test results once the resource owner is "of age".

Thinking modestly once again 🙂, we do have to prepare for the technical layer to fail gracefully in a privacy-protective fashion as the other layers vary.

The technical layer (modulo our question about a client turning around and becoming a downstream RS?) stays the same no matter what happens at the other layers around it. The legal capacity picture around it is *really interesting:*

- When you're young, you time-in to legal capacity
- When you're old, you time-out of it

- At some times in life, you might intermittently go into and out of it
    - Drunk or otherwise impaired (meds/drugs, distracted, tired)
    - Coerced (gun to your head)
    - Dementia but functioning
    - In Ontario law, there is a notion of "assumed implied consent" that lets someone take over for you if you're bleeding out on the ER table – a kind of break-glass right

How can we model this? Is it all down to revocation of the PAT? UMA doesn't have a notion of multiple resource owners, which presents a bit of a challenge.

John shared this blog post about privacy and fitness devices.

We had been thinking in terms of Ts & Cs that are relatively static for the (mostly) pairwise relationships between parties. Where could "smart contracts" introduce dynamism/configurability? The obvious place this had come in was in trust elevation, which is a Grantor/Grantee relationship. Where else could it apply? See our "trust relationship" discussions in the April notes (start here and work back).

Counterintuitively, injecting more democracy into previous feudal systems has the effect of creating more dynamic stability. That's what we're talking about here. (Not to put too fine a point on it. 🙂)

AI: Scott D: Take Eve's slides containing four use cases and "translate the technical stability into a persuadable view of stability and value in other regimes" for use in the primer.

2016-09-09

- Working session on User-Managed Access (UMA) in Contractual and Regulatory Contexts

Attending: Eve, Ann, Kathleen, John W, Adrian, Scott D

We discussed Eve's new slide describing "key benefits to users" at a high level as a potential way of fleshing out the next subsection of the primer.

- Not just opt-in or opt-out when asked
    - Sharing, unsharing, and editing of sharing preferences allowed at any time, without external influence
- Possible to offer a service that centralizes sharing preference management across data services for user convenience
    - The central service doesn't see any of the data
    - It acts on the user's policy instructions when others attempt access to data services
- The user can choose to share whatever "grain" of access each data service offers
    - Such as read vs. write, or weight vs. fat mass

We've now added and wordsmithed this.

Instead of the spiral or the simplified spiral in the primer, John suggests a very friendly version of the "three phases" paradigm (as suggested by Adrian) that's already explained in the spec. He will sketch what's he's thinking of, and we'll ask Domenico to turn it into something beautiful!

2016-09-02

- Working session on User-Managed Access (UMA) in Contractual and Regulatory Contexts
    - Eve will try to press ahead with lots of editing AIs prior to the call
    - Adrian and Kathleen have sent various suggestions in list/private email in the last month we should review

Attending: Eve, Kathleen, Ann, John W, Mary, Jim

We did a ton of work in the document.

If you haven't seen it, the latest version of the slides with the "legal use cases" is here. Please feel free to share it.

See also Jim's CommonAccord capture of the GDPR.

2016-08-19

- Working session on User-Managed Access (UMA) in Contractual and Regulatory Contexts
    - We can make progress with the current scenario

Attending: Eve, Kathleen, Mary, Jeffrey, Jim, John W, Ann, Mark, Colin L.

NOTE: No meeting next week!

The "UMA technical" agenda is pressing ahead. Eve summarized our UMA legal status for the LC as "fits and starts but progressing" recently. The idea of changing or providing multiple scenarios is fine; doing that after we press ahead on the doc seems like it won't block our mapping work.

Jeffrey's Information Governance fact pattern includes Home Trunk Industries, which decided to be an aggregator. The challenge in applying this directly to the UMA scenario is that an aggregator usually gets full control of all that data. Are some "orchestrator" companies (service operators) able to use new consent strategies that defer to Alice the ability to control her data and even access to her physical stuff? Eve just had a Twitter conversation about car /plane co-leasing with maintenance taken care of through UMA, and we discussed drone and lawnmower control through smart contracts. John puts it this way: HTI would own the data (it's fungible), but Alice would own the information (it represents her).

We made lots of edits to the Introduction and nearly finished it.

2016-08-05

- Review User-Managed Access (UMA) in Contractual and Regulatory Contexts
  - Nailing down the right 1-2 paragraph scenario for use throughout the document (send your alternatives before the meeting)

Attending: Eve, Kathleen, Andrew H, Adrian, Mary, Jim, John

NOTE: No meeting next week!

Regarding finding the right scenario: Smart home scenarios are really just as important due to sensitivity as health scenarios. In health research, it looks like the first BSC use case is going to get more complex still because real-life scenarios are getting blocked due to patients' reluctance to consent without a way to be notified of additional researchers' requests to access. Jim notes that "broad" consent is the key problem – anticipating future needs. And speaking of this, Adrian dislikes "informed" consent, because people don't know what's going to be done with it later. Transparency in being kept aware of status and where else the data is traveling is a key part of resource owner empowerment.

Our discussion this time is ranging all over, but it all seems related to our hoped-for Regulatory section.

Maybe we need a discussion in the Regulatory section about what "consent" maps to in UMA's "authorization" concept. There is a notion of authorization as a thing that can be dynamic – an RO can grant authorization and later revoke it (by an interaction with the AS that is actually outside the protocol but in scope for our model text), at a grain that is finer than Ts & Cs, and empowered by that grain, by the option to "Share" at will, and by the option to "un-share" any portion of the resource at will. John likes the phrase "Dynamic Consent"... It may be taken already. See this Nature article.

Hazard@All: could have an obligation on the data user to destroy. Like my passport number for a hotel stay.Hazard@All: but once the other guy has it, we need to rely on some other method to have them delete it.Hazard@All: either contract or escrowHazard@All: that is a matter of legal verbiage, I think Hazard@All: the lawyers phrase it however seems to workHazard@All: and the vocabulary will change based on dozens of factors, including the common vocabulary of the business sector, jurisdiction and languageHazard@All: notice can be specified however you want - right?Hazard@All: you can use my stuff but you agree to text me whenever you do.

Would we have to deal with privacy/data protection regulations as property-based vs. rights-based in our model text? We may very well, but it needs more discussion. Could Jeffrey apply his model?

AI: Eve: Ask if Jeffrey would be willing to adapt his Information Governance school assignment into a small scenario for the primer, adding Alice to it, as an alternative scenario for consideration.

2016-07-28

- Review User-Managed Access (UMA) in Contractual and Regulatory Contexts

Attending: Eve, Scott, Adrian, Jeff S, Jeffrey R, John, Kathleen, Jim

The 50K-foot view: We are tardy on our original model definitions and clauses schedule, but our "primer" deliverable is intended (ultimately) to drive clarity and force around the model text deliverable.

Do we want to enlarge our scope to encompass "enterprise UMA"/federated authorization use cases? Maybe we give it a nod, but don't address it fully. We may want a separate, similar paper on the "Legal Person" opportunities.

Our history with BLT in the primer: It didn't survive our analysis of doc structure. Scott's history with the invention of BLT: It was meant to be extractive: breakdown and then buildup. It might be the case that the duties within our obligations in the model clauses have "BLT metadata" somehow, if it's determined that this is useful, but they seem not to be a useful organizing principle in the doc.

In the intro, we want to get from a sharp Problem to a Solution and then to an enticing How we're going to help them solve it. UMA can help soft, fuzzy circumstances become harder and more measurable.

What scenario should we pick? Should it have "health" in it? We always debate this. Should we go with "smart home" and non-health, perhaps? Or "pure Google Apps-like"? We could tweak it to have smart home for a disabled person to allow them to live at home vs. other circumstances. That introduces an additional layer of regulations, however. The appeal of health is that the data is so intimate and sensitive. A baby monitor, or some other smart home devices, might have the same characteristics, though. IoT seems definitely popular.

Homework for next time: Please send your own ideal 1-2 paragraph scenario to the list so we can nail this down next week and power through the next section.

2016-07-22

- Review UMA Roles and Responsibilities Primer

Attending: Eve, John, Sal, Ann, Kathleen, Colin, Adrian, Jon, Scott (Maciej regrets)

We reviewed the primer and made changes dynamically.

We discussed whether "Alice and Bob" are appropriate as standard RO and RqP names, given their history as equal peers in PKI, vs. their asymmetrical nature in UMA. UMA's historical goal is to empower the RO role in an ecosystem fashion a la Eve's 2008 diagram from the Digital Contracts MIT event, even if instances of a running UMA protocol are indeed asymmetrical, since the RO has an "agent" – the AS – working on her behalf whereas the RqP doesn't. Of course, there's an admittedly asymmetrical part of the UMA protocol called trust elevation where the RqP may have to supply or direct services to supply information about themselves, which could be UMA-protected, so the RqP could have an "agent" acting on their behalf too. In short, while they're not peers at an "individual instance of the protocol" technical level, it's possible that they're similarly protected at the technical, contractual, and regulatory layers.

Self-regulatory structures are the common theme coming out of the list of "Additional Discussion Topics", which we spent a lot of time on. Some of those structures will be shaped by supply chain necessities – see, for example, how UPS and FedEx end up with similar contractual outputs. We think that once we have our "Tech/Contract/Reg" framework, we could spin out white papers on each of the additional discussion topics, and likely liaise more effectively with the IDoT, IRM, and BSC efforts in Kantara. (Be sure to see the notes so far from the BSC group.)

We reviewed the scenario in the Tech section. Scott asked: Do we want to draw the equivalence between the federated identity progress and federated authorization? Eve thinks this depends on whether we're writing for an audience that understands federated identity trust frameworks. Adrian goes by the "There's only one Alice" mantra – the AS shouldn't be "domain" (sector) specific. This is the personal AS viewpoint.

2016-07-01

- Review document deliverables
    - UMA Legal Primer – intro and first section have been fleshed out and there are some comments/questions that need answering
    - UMA Legal Use Cases

Attending: Eve, Kathleen, Colin, Jim (regrets: Mark, John, Mary)

We took all our notes in the doc. Wonderful progress. Thanks, all! We will try and make progress in the doc for next time!

2016-06-24

- Review document deliverables
    - UMA Legal Primer
    - UMA Legal Use Cases

Attending: Eve, Kathleen, Ann, John W, Adrian, Scott, Paul, Mark

We started reviewing the primer. If we can't make it all fit into 2-3 pages as we first hoped, let's try and use a "progressive disclosure" approach, so that the first part boils down the rest of the paper in a very short space.

NOTE: All but Eve (and John): Eve will "take the pen" over the weekend in the doc. If you want to do stuff in the doc till the next meeting, please do so in Suggest mode!

Adrian introduced the notion of the AS as the only safe way to handle bots (such as Siri) that are the alternative to the explosion of apps on mobile devices (see the de-app-ification trend). So the "agent" theme as exemplified by a *personal* AS could be an exploratory theme in the doc.

How deep to go on privacy compliance? Scott suggests that UMA makes the underlying systems reliable and predictable, which provides a basis for trust and enables observability of regular actions. The systems are tuned to local expectations. Things are "done to spec". Adrian observes that UMA could be said to be a "core component of privacy engineering". Scott also reframes as "operational privacy". Don't forget about Privacy by Design – here is the (old by now?) paper on Privacy by Design implications of UMA.

Why could an XXX love it?

CPO:

- Standards make for cheaper solutions for compliance. Average US privacy budget (PSR '15): $300K.
- Emerging ecology of user control standards gives an alternative to contracts of adhesion.

Data subject:

- There's power in having more "consent tech" solutions available and deployed on your behalf to choose from.

2016-06-03

- NOTE: No subgroup telecon next week
- Privacy@Scale report
- Status of outlining/editing of newly planned document deliverables
    - UMA Legal Primer
    - UMA Legal Use Cases

Attending: Eve, Paul, John W, Mary, Ann, Kathleen, Adrian, Mark

Privacy@Scale and other event reports: John and Eve both attended on Monday of this week; it was the 2nd annual, and their first. Facebook hosts it. It was in DC this year. John was very vocal. You can find two of the three report deliverables commissioned by Facebook from Ctrl-Shift at this site.

The third report is coming soon. Eve spoke at FB's behest on the roundtable process that fed into those reports, and mentioned UMA and Consent Receipts.

John points out Omri Ben-Shahar's research (he wrote a book called More Than You Wanted to Know: The Failure of Mandated Disclosure) purporting to show that privacy notices are pointless (he was doing a "point-counterpoint" talk with Lori Cranor). Mark contrasts this (?) with the Project IF work on "data licenses", where it seems to matter.

On Tue-Wed of this week, Adrian attended the NIST Named Data Networking meeting. This work has significant privacy implications. Relative to UMA, there was a mention of longitudinal health records. But he doesn't expect movement on this soon.

PPR has its annual Health Privacy Summit. Remember: It's free and streamed. Contact Adrian to get an invite!

New UMA Legal document deliverables: John did a draft outline of an "UMA Legal Primer". Eve has her slide deck as a start for something like an "UMA Legal Use Cases" document. What should the real names of these documents mean? "Legal" comes after something like "Business", "Strategic", "Real-Life", etc.

The use cases are ideally tackled in a fashion that doesn't introduce UMA terminology, at least at first. They can introduce UMA terminology as a mapping exercise in a second chapter or something, or even point to the other document.

What real-life digital resources should we use for examples in the use cases? Health data is a very real example, but how much should we use this? There's a sense that we should either spread the examples around, or be cautious/aware about using health examples because of its specialty characteristics.

We could introduce RACI mappings once we come out of the "UMA technical" term explanations into the "data *" mappings. Kathleen notes that when we have a 1yo (actually a minor who is legally incompetent but old enough to be aware) Johnny, he is legally incompetent to consent and thus may give assent vs. consent; this maps to "Resource Subject that is not at a Responsible level" or something. IOW, RACI is an imperfect system to map to (as we concluded previously). However, the point is that RACI maps well to "data *". (John writes this part. :-D)

Further, mapping UMA-technical roles to data-* roles leaves a hole: the AS. It's not a data-anything. This is the innovation of UMA.

The contract roles and relationships is the third plane. We should treat it on its own.

The goal of the UMA Model Clauses section would be melding the three dimensions all together. We should take a lightweight approach, point to a couple of illustrative clause samples, and in essence serve as clause documentation. We might need a subsection somewhere in this doc (not as a separate dimension/plane) for the UMA-legal definitions, to support the clause explanations.

Eve (John?) will import this into a GDoc for collaborative editing.

2016-05-27

- Review Digital Contracts event
    - DG
    - CommonAccord directions
    - Legal POC appetite
    - ...
- IDE directions
- Model definitions: next steps

Attending: Eve, John W, Adrian, Ann, Jon N (Andrew regrets)

Adrian is on the FHIR group, along with Kathleen. The conversation at that table around "policies on the wire" is very live there. The motivation to source policies from multiple places in the FHIR conversation seems to be internal to a single domain.

Given our capturing of notes last week, we'd like to write a short document for a nontechnical, legally inclined, even regulatorily inclined audience. John has stuck his hand up to outline such a doc, and he and Eve can review the results together at Privacy@Scale (run by FB) next Tuesday. (Microsoft is running a similar event the next day, NIST is running a Named Data Networking event on May 31-Jun 1, and there's a Health Privacy Summit at Georgetown Law June 7-8, and it will be streamed.)

Looking at it from a business/strategic point of view, the benefits are privacy-preserving because no authorization/permission policy transmission is needed, and the AS is not a data controller, and individual/patient/consumer/citizen-centric because the AS exists to enable that party's wishes (their "agent" in some fashion). Looking at it from a legal point of view, we are working on a system of model text to protect the interests of all the parties engaging with each other.

(As we said last week, if an institution running that domain wants to federate policies in some fashion a la the methods that XACML makes available, they're free to, behind the "UMA façade" of the standardized APIs of AS. If it turns out to be valuable to develop a companion technical paper that answers questions that technical people have about the architecture that supports our contentions, we could write that afterwards.)

It seems like we really need a proper use case document now too, identifying why and how the parties might be equivalent or not equivalent (note that in UC4, "offlice Alice/gov agency", the Grantor happens to be the ASO, but in UC1-3, it's not). We need more use cases showing how the Grantor could equal the Grantee, the Grantee could not equal the human end user that is acting on behalf of the Grantee (like Dr. Bob working for the hospital that wants access), etc. Once we have a full, non-repeating set of use cases, we'll know we're close to being done with the model term definitions. 🙂 Eve's interested to coauthor; John W's tentatively interested to coauthor; Jon N's interested to review.

There's a new Kantara Blockchain and Smart Contracts Discussion Group – feel free to join!

Note that the actual original EU Model Clauses are under threat. But we seem to be happy with the lowercase phrase for our work, still.

We should have a CHEDDAR rundown on a future call.

2016-05-20

- Review IIW and EIC discussions – progress around RSO-ASO as an "agent"
- Upcoming Digital Contracts conference with opportunity to gather requirements for CommonAccord IDE and feedback on our work
- Reality check on model clause process – definitions seem a lot easier, vs. clauses without a "live" agreement to work towards
- Discuss: What about the Grantor keeping policies secret?

Attending: Eve, Adrian, John W, Ann, Mark, Kathleen, Jim, Colin

At IIW, we reviewed the model definition work to date, and discussed the use cases for splitting the Grantor and Resource Subject. We discussed how the ASO is an "agent" of the RO/Grantor/Resource/Subject. In the "J-LINC - Signed Contracts on a Blockchain" session, we discussed the three different kinds of delegation. See yesterday's WG notes for more on this: we said there are different kinds of delegation:

- RO-to-RqP (delegating access to resources – the problem that the UMA protocol "solves")
- RO-to-AS (the architecture that the UMA protocol invented to solve the problem, which makes the ASO an "agent" for executing Alice's wishes)
- Resource Subject-to-Grantor (offline wrt the protocol, but an important business/legal relationship where the Grantor assists the Resource Subject in controlling access by RqPs)
- Authority-Override (offline wrt the protocol, but reflects jurisdictional overrides to ASO or RSO or CO actions and conditions in the business/legal space)

In the regulations generally, there's a notion of a data subject, a data controller, and a data processor. But there's not a notion of *another* player: a service, which UMA calls an Authorization Server – *and* an operator of that service, an Authorization Service Operator -- that executes the wishes of the data subject. And at the legal level, UMA has observed that there may also be a Grantor who assists a Resource Subject in controlling access.

When we say "offline wrt the protocol", it's offline wrt UMA itself, but there's the possibility of a notification pattern that could be implemented in a technical solution. Right now that's not quite on the table for UMA, though the CISWG is discussing this somewhat. There are six ways in GDPR that an entity could collect personal information, and only one is through consent. Proof of consent is required, though not to be presented to the individual – rather by the org to authorities. Other open principles, however, do require presenting such proof to the individual. Canada in some cases does require proof to be presented to the individual depending on information sensitivity.

Let's stick strictly to a RACI model (responsible - accountable - consulted - informed). Is it possible to assign static RACI roles to ASO, RsS, and G, or not? If "little 1yo Johnny" is a RsS but not a G, he can't be responsible, only informed, right? Every consent has to be scoped appropriately, and it has to be dynamic.

The RSO is most clearly a data controller. The biggest conundrum wrt regulations such as the GDPR is that the ASO isn't a "data *anything*". It holds policies, which are metadata, which don't necessary have to travel anywhere, and – in UMA flows, anyway – don't travel. Kathleen has brought up previously how you can have "federated" flows behind the UMA scenes where policies can be sourced from multiple places to make an authorization decision at the UMA AS, and technically, the UMA AS could have gotten the token that it handed to the RS from another place entirely. In other words, the job that the AS does could be entirely virtual. In the main, we might guess that an AS is "monolithic", but what if the job that an AS does is through an API facade, and is really done by a bunch of microservices run by lots of different companies on the back end?

Netting this out: UMA uses the OAuth architecture, which only puts calculated entitlements on the wire, not policies, which seems to be a privacy benefit (vs. an XACML-type architecture, which designed a way for policies to go onto the wire, and also requires the "RS" – the PEP – to ask the "AS" – the PDP – for the policy decision, vs. making the "client" – the application ask for it).

So where we are is that the AS has a strength in not being a DS/DC/DP. It is a helper service; an enabler of the Grantor; a technical device for connecting services and data stores; a data sharing and control mechanism available to the data subject. It's a tool to enable Alice to control who has access to her data. The AS is the linchpin of the UMA protocol; the service at the center of the marvelous spiral.

How can we make progress on real live model clauses for real live agreements? Eve might have someone interested in a "legal POC". Let's ask people next week at the Digital Contracts event too. Mark has someone in mind too. Adrian suggests that a MITREid Connect/HIPAA context project could be a good case. Have contacts reach out and we can discuss starting "legal POCs"!

2016-04-29

- Accountability (data controller) vs. responsibility (data processor) in relation to the data subject (this topic came up at our UMA Legal IIW session)
- Terminology: the next generation – can we tackle Grantee subtleties next?
- What is the right structure for UMA Legal work and review? (this topic came up at our UMA Legal IIW session)

Attending: Eve, John W, Adrian, Jon N, Mary, Dazza, Ann

RACI charts map the responsible/accountable/consulted/informed split. We discussed at IIW22 how the accountable body, the data controller, would typically outsource to a third party, the data processor, some work. It's optional to outsource. Would could it possibly mean if Alice really runs her own AS (she's the ASO) and her own RS (she's the RSO)? Could the AS be viewed as a DP in any sense? We recognize that conceiving of the DS as a DC is really novel. The IIW microservices conversation is relevant to this because micro-level service conversations that go cross-domain don't yet have a "legal model" of personal data flow.

Each RS seems naturally to be a DC. Okay, but what about the case where you have APIs that allow clients to put data back in to the service? E.g., if the API presents a POST-based call that lets a multitude of clients send various kinds of PII-bearing data to the service for safekeeping? A cloud file system is a great example; there are lots of third-party GDrive clients. Maybe there's some split of DC and DP roles among resource servers and clients for any one API, determinable by their contractual trust relationships. We didn't even list **RSO - Client Operator** as one of our trust relationships below, because so it's clearly out of our scope.

An important note: Alice's interactions (in whatever technical or legal role) with the AS and RS are actually out of band; there are no UMA messaging arrows there. There's only "contractual" (or role?) and "regulatory" reality.

John W is suggesting that we have two diagrams of UC1 (etc.), the current one and a new one that carries an analysis of how the DS, DC, and DP roles might map onto our roles.

Jim and Thomas have an opportunity, if they wish, to bring it up at the GLTL MIT event they're attending next week. The exercise would be to map the flows of information represented in the use cases to the delegation of authority and responsibility represented by the model of authority and responsibility in the GDPR.

Is the AS even anything in this system? Maybe nothing; it's novel. Jon N believes it doesn't add anything and could make things work ("worse"?). UMA brings something new, and the problem is of consent. The AS, in GDPR terms, must stand in for the DS, and the critical problem is that it may be difficult to make the case for this important role given the regulatory regime that never imagined this role. The AS needs to be the DS's "agent"! How can we achieve this? Let's seriously consider outreach on this score as part of our remit.

Watch out for missing meetings in our schedule.

2016-04-15

- Grantee and Resource definitions; how to handle role stacking (Grantor can be Resource Subject, RSO can be ASO, etc.) – handle in a generic way? (see this thread)
- (Based on input from Scott D.) RSO - ASO model clauses (original proposal):
    - How to parameterize to enable different legal models for handling the granting of "access"?
    - How to identify/parameterize to enable exceptions to normal access/non-access? (see decision tree)
    - How to handle notification requirements?

Attending: Eve, Ann, Adrian, Jim, Kathleen, Mark

Adrian's idea for an IIW session: UMA and GDPR.

Do our model clauses actually want to be broad enough to serve the use cases for things like the enterprise as a resource owner/Grantor? So far, we do have a definition of Grantor etc. that admits a Legal Person (e.g. an enterprise or government agency or whatever).

HL7's definition of grantor: "An individual who agrees to confer certain rights or authority to a grantee." A grantee is: An individual who accepts certain rights or authority from a grantor." The various ISO and EU standards discuss collection, access, use, and disclosure. Unfortunately, they didn't think of the case where API access is about other HTTP verbs than GET 🙂, where the client might be putting data for which it's authoritative back into the server, or whatever; in these cases, collection, use, and disclosure could sort of be reversed. This is why "access" so often carries the burden for the other three.

Jim believes: "A goal might be to create text handles for each of the foreseeable circumstances. So, a complete vocabulary seems better to me." So shall we define the full set, then? We would theoretically know when to use which phrase (potential vs. real) in each clause. After all, the "potential Grantee" doesn't have consent, right? Why conflate?

What about carefully staying away from applying a specific legal theory to underlie the agreement (e.g. license)? We've already put out of scope the different agreement life cycle templates; Jim has done a ton of work on this in CommonAccord. (Kathleen is interested to learn about this so she can apply them in status definitions in the HL7 contract work!) Essentially, we think staying away from picking normative legal contexts, including workflows, is still out of scope, so wherever this might "intrude" on our clause text, we'll need to figure out how to parameterize the text.

Next time without fail: Let's focus on the #RSctrl clauses needed (which may be part of the RSO - ASO trust relationship and also part of the RSO - Grantor (Potential Grantor?) relationship). See the Jan 15 notes for four use cases, and also realize that we have to take into account a fifth use case of "requested access outside of normal UMA mechanisms".

We discussed the GA4GH work on "ADA-M"; they're developing what look like model clauses for the sharing of genome data.

2016-04-08

- Our budget request was accepted! Let's work out a timeline
- Work on RSO - ASO model clauses more – can we get to "beta" quality today?
- How are invitations to our April 15 meeting going?

Attending: Eve, Andrew, Ann, Adrian, Mark, Kathleen

Budget request accepted

The Kantara board has asked us for a cash flow timeline. Eve's estimate was that we could develop our requirements by early Q3, say no later than July, and identify a contractor by early-mid Q3, and ideally conclude the project by the end of Q3 (Sep). This gives us a forcing function to have a credible set of beta model text to work with by Q2. We seem to be able to accept this, even though there are some question marks.

AI: Eve: Reach out to our community of "legal eagle" participants to wake up active participation once again.

Realistic first beta review timelines

Is it realistic to invite anyone to our Apr 15 call? Are we ready with text? There are two parts to this: Can external parties attend, and do we have enough text/materials?

Adrian proposes that Kathleen is on the critical path here. Kathleen just wants to run a FHIR pilot based on our clauses! And doesn't see Andrew's use case as incompatible. What's most practical? Kathleen suggests that getting to the point of real clauses is really the most practical. Let's work hard to continue putting our model clauses together, and reassess just before the IIW timeframe as well as think about a June forum for external review.

Model definition and clause work

The FHIR connectathon is in May. She would like to "XMLize" the clauses for that target. What actual language is being executed here (since XML is just the "punctuation", if you will), if FHIR has a place in the code for a Grantor etc.? It sounds like they are using UMA authorization servers to consume standardized policy expressions in XML (if not necessarily XACML) format, and this would map to the Grantor's wishes for sharing. This could come down to a particular set of trust relationships, and the question of whether it splits into "active" policies vs. things like consent receipts.

Our Grantor - RqP (Grantee?) trust relationship already mentions consent receipts, and we've already discussed how CRs need to be part of a workflow that shows the Grantor what it is they're consenting to (or proactively sharing/delegating or whatever, i.e. setting a policy/sharing preference), so we could already be in the ballpark here. FHIR is using its audit logging framework for this workflow.

Adrian had made a generic "NPE Release of Information Form 1" sample form a few months ago, and Kathleen observes that with a few additions this could be something like what FHIR needs. Adrian is intending this to be used by the RS/RSO for setting up phase 1 (that is perhaps a rathole for our purposes for today's agenda; Eve suggests that Adrian and Kathleen discuss offline as appropriate).

Eve shared a [slide deck] that runs through various design patterns for splitting the technical resource subject and legal Grantee roles. This deck now has some explanation on the first slide about how the technical UMA flow can stay the same, while the trust relationships and legal/contractual backing might vary significantly.

2016-04-01

- Work on RSO - ASO model clauses
- Decide external review strategy

Attending: Eve, Andrew, Adrian, Kathleen, Thomas, Jim

We discussed the candidate model clauses for RSO - ASO.

What is a resource set description vs. a scope description vs. a resource vs. a scope etc.? Should we be really specific, the way a technical spec – or legal document – would expect, and define every little term? A resource set description being registered is a set of metadata about an *actual* digital data resource of Alice's that is being put under the AS's protection (in other words, the data itself is not uploaded), and the scope description(s) that are part of that metadata are descriptors of the *potential* extents of access that are possible to be performed over that resource – not *actual* entitlements that some requesting party has "won". One metaphor is that a resource set is a "noun" and a scope is a "verb" (among some set of verbs) that is possible as an action on that noun. Another metaphor is that the scopes are the definition of the contract, or interface, for that resource. In effect, the scopes define the "API" for the resource. Note that since RS gets to choose how to split up its design of resources and scopes, and the design line between them is moveable.

subj - verb - object (crude representation of an authorization policy)

RO - RSO - RSO (who's authoritative for each piece)

ASO (authoritative for executing the whole thing)

Adrian's concept of "I split, you choose" means that the RSO gets to be authoritative for that digital data resources and the APIs presented for exposing them. UMA has innovated around enabling the RS outsourcing the function of letting the RO be authoritative for associating the "verb and object" part of a policy with the desired subject, and then letting the AS execute the RO's wishes. This is through a combination of resource set registration and then the rest of the UMA flow.

Do we have a chicken/egg problem in terms of the relationships? What if there are two independent relationship streams leading up to the three-way PAT relationship?

- Trust relationship 1: Grantor - RSO (no technical artifact)
- Trust relationship 2: Grantor - ASO (no technical artifact)
- Trust relationship 3: RSO - ASO (technical artifact is OAuth client credentials for the RS)
- Trust relationship 4, dependent on the previous three: Grantor - RSO - ASO (technical artifact is a PAT, carrying user consent)

(The numbers where there are no preconditions are arbitrary – they're just for easy reference.) There may be cases where the conditions are variable. Eve was concerned about our two-party "atomic" approach previously because a PAT inherently has three entities involved, and it naturally maps to three parties; likewise, an AAT. Her suspicion is that some of our duties may have to "go there" in mentioning three parties, and that this isn't a variable thing: the clauses will have to do that.

Adrian asks: Is it valuable for trust relationships 3 and 4 to have different model clauses? In other words, would the RSO and ASO ever want to contract with each other separate from the context of any specific individual? Eve made the case that some sort of blanket protection against liability, in both directions, makes sense regardless of the particular resource owner. Kathleen mentions market forces as pressuring ASOs and RSOs variously on this point. Also, regarding #ROctrl, that's going to have to be fought out in trust relationship #1, not trust relationship #3; the ASO is more likely going to want to say "Listen, I'm going to put stuff in the token, and if you do something else, I wash my hands of you."

Regarding external review, Andrew would like to point some DIACC members to our model definitions.

AI: Everyone: Reach out to external reviewers and offer to hold a live review session on Apr 15 and see who can attend. Eve to send a reminder of what this is all about in email.

Let's also plan a live review at IIW, by which time we'll have more model text.

2016-03-25

- Work on RSO - ASO model clauses
- Choose first trust relationship to work on for model clauses
- Identify any model definition areas that need work given our choice
- Decide/confirm our plan for first round of external review

Attending: Eve, Paul, Kathleen, Adrian, Jon N, Ann, Jim

We reviewed our [model text] with [model definitions], with recent changes.

What's the deal with the "D word" – delegation? If we qualify "delegation of *what*", then we can observe that there are several types, and we can map them to different trust relationships we have identified:

- Resource Subject - Grantor (not in scope): This enables a "resource owner" to find someone else to actively manage resource access for them at an AS.
- Grantor - ASO: This enables an always-on AS to be available to execute the wishes (policy conditions) of the Grantor.
- RqP - Client Operator (not in scope; also not previously discussed as a model definition...)
- Grantor - RqP

Discussion about Grantor - RqP: Adrian doesn't think this is an appropriate use of the word "delegation". Eve believes "that train has sailed". 🙂 See, e.g., the NZ POC.

Regarding Adrian likes the distinction of "If you have a privacy policy, then there's an Operator involved; if you don't have a privacy policy, then you own it /wrote it." (It's personal to you and there's a 1:1 relationship between you and the software/service.)

We have a new party with responsibilities now in scope! It's the Client Operator.

Also, we're likely going to have to change references in the bodies of various model clauses from the Grantor to the Resource Subject, since the duty will be owed to the Resource Subject instead in a lot of cases.

Which trust relationship should we dig into first? For reasons of our #RSctrl/#ROctrl use case tension, and because it's the most "upstream" (no dependencies) trust relationship, let's start from RSO - ASO (see the clauses that start ASO... here and the clause that starts with RSO here – however, these are old and almost certainly don't reflect our latest discussion).

There's a use case that could help us interestingly when we get to working on Resource Subject- and Grantor-involved model clauses, but will need to take into account some out-of-band machinations because it engages precisely the Resource Subject - Grantor relationship: a "digital death" use case. Alice (as her own Grantor) shares some digital data resources with Bob. She wants to arrange for her lawyer to become her Grantor when she dies, so that Bob can continue to get access to those resources. In practice, because of what are known (really) as "time-to-live" strategies around credentials, tokens, and permissions inside tokens, the access that Bob enjoys to those resources is certain to get cut off at some point after Alice dies. Using a combination of business, legal, and (non-UMA – including digital contract) technical approaches, Alice's lawyer can replace Alice as the owner of those resources and ensure through the AS that Bob's access can continue or be restarted.

AI: Jim and Eve: Extract the two-party reworkings of the old model clauses for at least RSO - ASO and prepare for group review next week.

AI: Everyone: Think about and approach candidate external reviewers by next week.

2016-03-18

- External review and reviewers of our work at the Q1 mark
- Action items for the model definition editors regarding last week's consensus definitions
- In-scope and out-of-scope trust relationships and their names (see this thread and below)
    - Disentangling meanings of "delegation" – do we need a group norm about terminology?
- How do we square the "Requesting Party"/Grantee/Potential Grantee circle?

Attending: Eve, Kathleen, John W, Sal, Ann (regrets from Adrian, Mark)

AI: Eve, Jim: Revise our CmA model definitions with new consensus definitions.

Eve's premise was that anything with an UMA technical artifact really has to be in scope, and a direct relationship with any service operator with an UMA-specific purpose for existing (namely the AS operator) seems plausible to be in scope.

Would it be useful to write an explanatory document for reviewers about these relationships?

Should we be naming the agreements at all? There will be multiple actual agreements, for example, at different phases, or for whatever reason. And (channeling Jim) there are literally different agreement/contract workflow phases such as offer, acceptance, etc., and different sides of the equation might be offering or whatever. And given that English is only our first target language and not our only one, perhaps it should be a principle of ours that we should name as few things as we can get away with, so as not to have "language skew". Since our ultimate deliverables are at the size of "clauses", we need not care (for purposes of discussing trust relationships) about the larger combinatorial things into which they're mixed. Jim has put together quite a few assembled agreements, contracts, offers, etc., and we might do the same with our clauses as exemplars, but in the main we expect they'll be non-normative. So we've saved the candidate names, but kept them non-normative.

We have established to our satisfaction that there is complete separation between the legal and technical work of UMA by demonstrating that a technically conforming deployment of UMA services could be out of compliance with the data protection laws of some jurisdiction. Our "Russian example" is apt here: A Russian Grantor cannot legally allow a non-Russian "Requesting Party" access to their own data.

The clauses need to set up the correct preconditions for a deployment of UMA that encourages a "race to the top" in a greater alignment of incentives among the parties and in behavior that is compliant to relevant nonfunctional requirements (such as regulations).

John notes: "Makes me think that we need something like: Authorizing Server Operator - Regulating Body: Allowed-To-Release"

Trust relationship	Dependencies on existing nonfunctional requirements	In scope for model text work?	(Candidate names of agreements formed out of our official clauses)	Technical artifact	Dependencies on existing technical artifacts	Discussion
Resource Subject - Grantor (if they are different)	n/a	No	n/a	n/a	n/a	
RSO - ASO	n/a	Yes	Protection API client agreement	OAuth client credentials for RS	n/a	
Grantor - ASO	n/a	Yes	Authorization	n/a	n/a	This has no technical artifact in UMA, and

			services agreement			therefore no exact or obvious moment of appearing "on stage" for our work. Does this have an impact on whether it is in scope based on the rationale above? We agreed there are definitely reasons to have UMA-specific clauses, looking at our draft ones.
Grantor - RSO	RSO - ASO	Yes	Resource protection agreement	PAT	OAuth client credentials for RS	It's awkward to have a two-party agreement resting on a PAT. There are plenty of agreements where there are multiple Persons in one of the roles, but we're not so sure about true three-role agreements.
Client Operator - ASO	n/a	Yes	Authorization API client agreement	OAuth client credentials for Client	n/a	We have not scrutinized the definition of Client Operator. Let's put a pin in that to be sure it's okay.
"RqP" - Client Operator	n/a	No	n/a	n/a	n/a	We need to work on the term for "RqP". Let's put a pin in that.
"RqP" - ASO	Client Operator - ASO	Yes	Resource authorization agreement	AAT	OAuth client credentials for Client	The same discussion about n-party agreements above applies here.
Grantor - "RqP"	Resource Subject - Grantor, RSO - ASO, Grantor - ASO, Grantor - RSO, Client Operator - ASO, "RqP" - Client Operator, "RqP" - ASO, Grantor - "RqP"	Yes?	Resource sharing agreement Resource access agreement	Trust elevation mechanism and specifically claims subprotocol	OAuth client credentials for RS, PAT, OAuth client credentials for Client, AAT	Could be in the form of a consent receipt, sent to a notification endpoint. Is this trust relationship too "dependent" to be in scope, or does the calculus work? More discussion is needed.

2016-03-11

- Do these terms work for us as model definitions?
    - Resource Subject: Person to whom a digital data resource relates
    - Grantor: Person who manages access to a resource, either as its Resource Subject or that Person's behalf
    - Grantee: Person seeking access to a resource
- How to address the subject/grantor split in the tough use cases
- Tackling model clauses for RO/RSO
    - Does the brute force method help?
    - How can we address #RSctrl and #ROctrl needs – can we actually enumerate the exceptions?

Attending: Eve, Kathleen, Scott, Adrian, John W, Jim, Mary, Andrew, Thomas, Mark, Maciej

(As an aside, the fact that we're using Thomas's paid join.me account and had trouble reusing it because multiple people on the call have the app is an interesting "property rights" use case that is relevant to our problem! There's the potential for misuse of the paid account and the company often stomps on simultaneous use.)

Under consideration:

Resource Subject: The Person to whom a digital data resource relates
Grantor: The Person who manages access to a digital data resource, either as its Resource Subject or on that Person's behalf
Grantee: The Person seeking access to a resource

Is Grantee too "final" (vs. Requesting Party) given that they haven't yet been given access/had trust elevated? Maybe we need to split the RqP language similarly into two, to capture the states of "authorization not granted" and "authorization granted". To take an extreme example, there could be a policy condition that Bob the RqP gets access to some resource only on even-number hours of the clock. In this case, the AS can manage the token time-to-live characteristics very dynamically, the RS can check all the token authorization data frequently, etc. – but Bob will go from being a "potential Grantee" to being an "actual Grantee" really frequently. So the former phase or state is the more persistent, and the latter is the more transient. Trust elevation and degradation in Bob can be quite dynamic. (Also note that the same is true for the client app Bob uses.)

Note that any trust relationship between the (requesting party bundle of concepts) and the Resource Server Operator is entirely invisible to UMA. According to our scoping of our work, it should likely be invisible to us too. ??

But the (requesting party bundle of concepts) and the Authorization Server Operator does have a trust relationship once the authorization API token is issued, so we need to take care of that set of obligations.

- RSO - ASO (client ID)
- RO (Subj/Grantor?) - RSO - ASO (PAT)
- C - ASO (client ID)
- RqP (trust elev/degr?: Grantee) - C - ASO (AAT)
- RqP (trust elev/degr?: Grantee) - RSO = no UMA relationship ???

The "Adrian clause" in the UMA Core spec is intended to reassure RSO's that they are free to elevate trust in the "requesting side" (that is, both the RqP and C) themselves. Note that the phrasing in the spec only allows further tightening; a clause a bit lower down makes clear that the intention is not for the RSO to enable *loosening* of access, although it's impossible in technical terms to prevent this. This is why we are working on how far we can take our model clauses to give RSOs the appropriate (and not *inappropriate*) flexibility to manage liability.

Eve suggests a group norm: Let's talk about "notice" and "consent" as consent life cycle words, and "notification" as what happens apart from that life cycle, so heads don't explode.

We agree on these model definitions for now:

Resource Subject: The Person to whom a digital data resource relates
Grantor: The Person who manages access to a digital data resource, either as its Resource Subject or on that Person's behalf

We have discovered an important "phase split" in considering "Grantee", and can't really use it outright as a wholesale replacement for Requesting Party.

For the world of devices, client apps, and so on, should we introduce "Agent"/"Entity"/something? Eve is wary of Agent from our legal discussions, but Kathleen could submit ideas to the list.

Eve reviewed her "brute force version" of the "RO-RSO" trust relationship. Note that the last bullet is the only one that represents a break-glass scenario; the party seeking access never established a trust relationship with the ASO, never got into the position of being a RqP/Grantee/etc., and is unknown (at least as far as UMA is concerned) to the RO.

Adrian has introduced the term "resource registration agreement". Is this a good name for the overall agreement representing this trust relationship? ("Resource protection agreement"?) How do we square the circle of having two "three-party" trust relationships? How should we name each of the trust relationships? Let's try and bat this around on the list before the next meeting.

2016-02-26

- Resource owners and those acting on their behalf: what language do we need to support the distinction? Any additional use cases? Any additional insight on technical design patterns?
    - We said we'd try to conclude this decision-making process by this week, and then move on to the decisions about RS actions in contradiction.

Attending: Eve, Adrian, Ann, Andrew, Mark, Sal, Kathleen, John

Kathleen had summarized it in our WG call yesterday that this is an "owner" and an "on behalf of" problem. Adrian asks: Is it clear that there need to be separate credentials for these two roles (something like subject and custodian)? Kathleen's work at HL7 includes three roles: the **subject** of consent directive (of governed information), the **grantor** (the consentor who grants access), and the **grantee** (the one to whom access is delegated). Eve is familiar with a project where UMA is being contemplated where subject and delegate are the terms used, where she suspects (**subject+grantor**) and **grantee** map perfectly onto Kathleen's terms, and she knows they map exactly onto UMA's **RO** and **RqP** terms.

Let's put the "third term" aside since it seems there's no controversy about the mappings (for the moment), and concentrate on the "split RO" nature of the use cases in front of us.

Do we have to distinguish "who logs in to the AS" vs. "who logs in to the RS only"? Adrian asks: In our use cases from last week, are there separate credentials? We need to know the exact UMA design pattern we're looking at to know the answer, because otherwise the specific identity ecosystem in place is unknown. Andrew suggests: Is the question really about whether the RSO cares who the ASO is? Let's say a hospital IT department is the RSO and the healthcare insurer (Eve adds: or EHR SaaS vendor or government agency or whatever...) is the ASO.

Going back to the key issue: Adrian presumes that the AS presents a certificate and the RS gets to look at that certificate and see if it trusts it. But Eve points out this is off-track for the subject of this call. Any AS cert is for the whole AS. Let's return to (**subject+grantor**).

Regarding the terms coming from HL7, FHIR is taking this nomenclature work even farther. Eve wonders: Could subject and grantor suffice as terms for us, even if the subject isn't even literally the data/PII subject? Like, e.g., what if the "subject" created the digital artwork that is being sold and the grantor is managing access to the painting on behalf of the subject because they're no longer legally competent to be party to contracts? In health, a donor or cadaver might not be donating data, but body parts! (?)

The current theory/hope is that **Subject** could be a good new term for the "true resource owner" who might not be legally competent or might not be in a position to manage resource access at the AS. (Note that "control" is the verb we use on the UMA marvelous spiral architecture diagram to denote what the RO does at the AS.) If we need to get more specific in any of our model clauses regarding resources that truly contain PII, perhaps we need to clarify in the definition that UMA allows for resources not to contain PII, and then define a more specific term, **Data Subject** or **PII Subject** or **Resource Subject** (as discussed in notes below), specifically for resources that do contain it. (Do we also have to define the term **Resource** and/or **Protected Resource** in our model definitions?)

More current theory/hope: **Grantor** could be a good new term for the "guardian" of a legally incompetent Subject, or designated "proxy"/"agent" for a legally competent Subject, who exclusively manages resource access at the AS on their behalf. [Added after the call:] Eve's theory about the UMA design pattern in this case is this: The Grantor is the *only one* that interacts with the AS (e.g., logs in there, creates policies, etc.) If the Subject is capable of managing resources in an online fashion ("manage" is the verb we use on the UMA marvelous spiral architecture diagram to denote what the RO does at the RS, but also the verb we use to denote what the RqP does at the client), then in fact what could happen is that the Grantor sets up policies for the Subject to become an RqP to their own resources, accessing them through RS-that-is-also-a-C, and requiring that the Subject use the same AS as the Grantor to enable full Grantor oversight capabilities. (This is a known limitation due to #wideeco realities.)

More current theory/hope: We could replace **Authorizing Party** with Subject and Grantor.

More current theory/hope: **Grantee** could be a good new term for Requesting Party.

AI: Kathleen: Share HL7-based definitions of subject and grantor (and grantee) with UMA WG so we can try out definitions for our own purposes. (ISO standard basis for the definition – what is the ISO number? -- can't be shared because they're not free standards.)

AI: Adrian: Work with Kathleen to capture the "joint control" of painting resource example and share with the list.

AI: Eve: Add GitHub issues for the new model definition ideas.

2016-02-19

- **Distinguishing resource subjects from resource owners:** Can we develop a cohesive system whereby "resource subjects" without legal capacity can have "authorized agents" acting on their behalf as "resource owners" as required in order to forge "resource registration agreements" for the purpose of UMA's phase 1 particulars? Do the use cases/design patterns provide any insights or challenges here?
    - See email thread

Attending: Eve, Andrew Hughes, Paul L, John, Ann, Adrian, Jon, Kathleen, Sal

Is it valuable to solve for a model where an agent can be working on behalf of resource owner vs. a resource owner?

The protean nature of the word "agent/agency" is troubling. Is there a good substitute word? If not, do we have to define *Agent for all of our terms in an UMA context? We did already have Requesting Party Agent. Perhaps, at best, we should define it operationally but stay away from legal subtleties.

We've said that resource owner = Authorizing Party. Does that work, or is it not equivalent? There are terminology questions and there are UMA architecture questions. Should we just wave away problems by making them equivalent?

- **1yo case:** What if the "resource owner" (let's say they're the "subject" of the data residing in the RS) is a one-year-old kid and their mom has to manage the resources by logging in to the RS? The child is not competent to contract, even if they're old enough to sign their name. **Guardian** is a good name for the latter.
- **12yo case:** What if the "resource owner" (let's say they're the "subject" of the data residing in the RS) is a 12-year-old kid and they're old enough to manage the resources by logging in to the RS themselves? The child is not competent to contract, even if they're old enough to manage resources online. How to architect the system and name the parties?
- **Intermittently competent adult case:** This is another tough one.
- **Competent adult case:** What if the "resource owner" (let's say again that they're the "subject" of the data residing in the RS) is actually competent to contract, but wants to have someone else manage resources for them online? There's a paper resource owner, but an online "executor" of resource management. What's a good term?
- **Digital death case:** After the "resource owner"...

Adrian's concern is what happens in phase 1. These use cases have different properties in that phase. Eventually (soon), we will be in a position to work on what's supposed to happen when RS's want to take an action in response to an access request that is in contradiction to the permissions contained in an RPT (requesting party token). First, we need to understand exactly "who" configured the AS to

By the way, all the same patterns could apply whether or not the resources contain PII or not. What if the resource owner created digital media that they want to sell? Is there a reason to distinguish in our terminology at all? What if a resource contains PII "in bulk" for many individuals (in directories or databases or other repositories)? This was the point of Adrian's example. Eve's point was, rather, that individuals might want to be protecting resources that don't contain PII. Okay, now we're on the same page! There are use cases for UMA that span "Alice" and "enterprise".

Let's try to conclude this decision-making process by next week, and then move on to the decisions about RS actions in contradiction.

2016-02-12

- If you have terminology comments, speak up; issue #240 is closed otherwise
- Update on overall roadmap: initial priorities settled; new #shoebox use case
- #RSctrl, #ROctrl, and #shoebox use case are related; let's pick our highest-priority use case and specific set of clauses to work on, and work on them right now! (current clauses) (current relevant issues)

Attending: Eve, John W, Kathleen, Scott, Adrian, Mary, Jim, Ann

Terminology

Shall we leave this till we have specific needs?

We agree that Individual is best, over something like Natural Person. (Meatspace person, bag of protoplasm, ugly bag of mostly water...)

What about "data subject" and "PII subject" (which terms are defined in different standards)? We could refer to them from our Individual definition, or put them in a concordance appendix that is separate. Whether it's "next to" or "separate from" the definition text – and either could work in CommonAccord, because it could be modularized nonetheless, the question is whether we in this WG want to be responsible for keeping that text up to date. We've discussed that we don't want to be responsible for text that isn't UMA-specific. Jim notes that "any law firm that wants to make themselves useful" could keep up a text module like this. We would divest ourselves of any responsibility at all for the freshness or staleness of such text if we simply advised our model-text users of the existence of that third-party model text, rather than pulling it into our own model definition. It appears there are at least two companies that provide this sort of concordance, but they are paid services. Nymity (Toronto) and Privacy Guidance (UK) charge $10K/yr (for content that goes well beyond just this bit).

"Model definition" is a good name for the specific kind of model text that contains a full term or abbreviation definition. If we want to put a "comment field" in any of these, what would be the purpose? It apparently shouldn't render in a "final" legal agreement; it should just be there to guide how a legal agreement builder should use the definition.

Not just about terminology but about "agency"

Is the RO always the data subject (etc.)? If elderly RO Alice gives access to her protected resources to adult child Bob to help her manage them, UMA makes him her RqP. How do we understand "agency" in this light? Note that UMA calls it "resource owner" only because OAuth does, and "ownership" is problematic as a term. "Resource controller" is more accurate, and it maps nicely to the old WG concept of the "split RO" problem. Note that we have introduced the Authorizing Party phrase on the legal side of our terminology to help get around the "owner" problem, which aligns nicely with the concept of who has the "authority" in each context.

Scott introduces the "**legal capacity**" phrase. Every jurisdiction has a different definition of the point at which a natural person has it, and different kinds of health data are regulated in each jurisdiction. Sometimes legal capacity is down to age, and sometimes it's determined by other criteria. John mentions

tests that a physician can perform to determine capacity. In the case of "aging-in" or certain other automatic types of capacity tests, if resources can have standard data taxonomies applied to them when RS's register the resources at AS's, then it might be possible for automatic policy changes to take place.

Kathleen mentions that HL7 wants to use our deliverables for FHIR contracts. This is for its work coming out in September. They'd be interested in both the model definitions and the model definitions, for "substitute decisionmakers" – Alice would be the grantor, and Bob would be the grantee. In some other contexts, Eve has heard terms like subjects and delegates, and citizens and proxies, etc.

AI: Adrian: Write up ROI form concerns wrt to UMA phase 1 implications for the group's consideration.

Next up

We need to ensure we have a very firm handle on the legal capacity in our terminology (and possibly clauses). And then we need to make good progress on our beta model text timeline, for all sorts of reasons (both IDE and customers).

2016-02-05

* Term definition work (see issue #240) – for reference, here are all the open #trust issues

Attending: Eve, Scott, Paul, Jim, John W, Adrian, Ann, Jon N, Mark, Kathleen

Regulation of consent

Should we be concerned about EU regulators "regulating away" the potential power of UMA around proactive enduring consents (e.g., policies that are indefinite until revoked) because on subsequent reference to them by an AS they're not "explicit"? Is there a way to influence the thinking of regulators and /or have two-way conversations with them so that implementations and deployments can give individuals the right buttons and knobs and UX's? It's certainly possible to make appointments with regulators. The US regime has a property basis, whereas the European basis is human rights. Revisiting one's consent in context (monitoring one's consent) would theoretically be a powerful way to exercise one's rights.

In the healthcare world, coming from paper vs. digital, notice was expensive. But in the digital world, notice is free at the margin.

Mark notes that: "We are working on a model practice papers to send to regulators wth the Kantara sponsored workshops." He will send a note to the list with more information as required. He thinks we don't think we need to worry about regulators regulating away consent directives, as consent is regulated by purpose and notice.

AI: John W: Find ways to reach out to regulators to start conversations.

Digital Contracts, Identities, and Blockchain - new event at MIT

This event has very limited seating and is invitation-only. If you want to attend, let Jim H know soonest! The notion is secure, DRY, peer-to-peer text objects handled as if they were software objects.

Term definitions

Our term definitions of record are here.

We can define whatever terms we want, but we don't want to "chase our tails". Agency (here meaning legal agency, the ability to take responsibility) is different from being a hunk of, say, software; software isn't a thinking thing. We're in the business of helping services that want to be, say, IdPs and merchant services and healthcare services *also* be UMA services to create the legal agreements they need, and since they'll be UMA clause novices, we're providing starter UMA clauses for them.

Note that an UMA authorization server operator is different from talking about a data processor or data controller. The the former term is an "UMA legal" term of art, and the latter are regulatory terms of art. Wherever we use the latter, we would have to refer to our source of the definition – ISO 29100 is what CIS refers to for consent receipts.

The pairs we currently have are:

* resource owner:Authorizing Party
* requesting party:Requesting Party
* authorization server:Authorization Server Operator
* resource server:Resource Server Operator

Where things break down:

A use case: Alice (Individual) wants dentist's office (Legal Person) to get access to her calendar for the purpose of scheduling a root canal. The dentist's office receptionist (Individual acting on behalf of the Legal Person, as an employee or contractor) tries to access the resource, using a client (Client).

Questions: Let's not overcook this. Think in terms (ha) of the next strawman iteration. Since no one expressed strong feelings about making any changes, we will leave the terms as they are till changes are forced on us.

1. Should we give a generic name to the clients, authorization servers, and such? Are they Non-Person Entities as a general category now? What would be the usefulness of this? **Don't add.**
2. We've called the receptionist in this use case a Requesting Party Agent to date. Is that right? What benefit does it confer? There is actually an UMA flow (or possibly several flow options) to which it correlates. **Eve suggests:** Keep for now and try to use it in practice to see if it works.
3. Does UMA want to trace duties to preserve the rights of the Individual? **No further discussion.**
4. Does Individual want to be Natural Person? **Don't change.**
5. Is the ability of the Individual to consent constrained in the circumstances of the transaction, e.g., by regulations? (Think in terms of term definitions for now. This is a much bigger question!) **Don't change anything.**.

It seems we have nothing more to discuss on issue #240 unless problems arise.

2016-01-29

- Review the new UMA Roadmap for 2016 page and the "#trust"-related use cases
- Pick off some of the new "#trust" GitHub issues to work on
- Take a look at the documents Scott David shared in email on Jan 15

Attending: Eve, Domenico, Jon, Jeff, Kathleen, John, Dazza

Regarding #RSctrl, John W notes that Russia is one case of a jurisdictional constraint where the cloud service has to be located in-country and that would have to override the RO's choice. Adrian has brought up healthcare-related constraints regarding delays in fulfilling the access request. In the case of cross-border transfer, EU adequacy rules for offshore transfer come into play. This is why cloud services have data sovereignty plays and data centers are coming up in All The Countries.

The actual UMA Core spec has a clause, which Eve has dubbed the "Adrian clause": UMA Core Sec 3.3.3: "The resource server MAY apply additional authorization controls when determining how to respond."

Essentially, at a T (technical) level, *IF* the AS and the RS are run by different operators, we have very little direct control over the RS going against the RO's wishes as expressed by the artifacts produced by the AS. This is why it's so important to look at the L (legal) levers we can control. The #RSctrl use case is on pretty firm ground when it comes to legal compliance. How firm ground is it on in cases outside compliance?

How far could we take mandating the usage of our clauses? It depends on the compliance situations and the jurisdictions. There are also technical solutions that can be layered on top of UMA, such as encryption. If regulation of encryption use is already present in an ecosystem, then very likely both the clauses and the complex technology can be mandated. (If it's an unregulated environment and/or some element of the ecosystem is commoditized and "free-wheeling", potential partners at the edge may walk away because encryption technologies are complex and add cost and friction.)

We looked at Scott D's mapping exercise. Kathleen knows of a similar mapping exercise having been done and will point us to it. How might we be able to leverage such work? Jon suggests that we can look at the breakdowns of common text vs. factored-out differences to help us structure the elements of our model text that, of necessity, get into jurisdictional specifics. We can hopefully structure our common vs. factored-out elements in the same way.

Eve suggests taking on the term definition work first, taking the many healthcare use cases as examples – and possibly writing the needed model clauses to motivate the right definitions. E.g., what if Alice needs to share access with a hospital, or hospital department, and Dr Bob gets access as an employee of that organization? There are questions around how a "Requesting Party Agent" would get defined, and possibly also "Client Operator". While health isn't the only set of use cases for this, it's 1) the hardest case, and 2) ready to try out our work!

It sounds like we need to have focused text-bashing working sessions; our calls are the times we have available to do this.

2016-01-15

- Review the emerging WG roadmap use cases
- Discuss the "RO access limit discretion" use case
- Discuss other open editorial issues for the candidate model text, as outlined in previous meeting notes below

Attending: Eve, Ann, Adrian, Jim, Jon, Mark, Sal, Thomas

What Eve calls the "Adrian clause" is this part of UMA Core Sec 3.3.3: "The resource server MAY apply additional authorization controls when determining how to respond."

What does the "RS authorization discretion" use case mean? Should the AS be made officially aware that the RS still didn't give access (similar to the "Shoebox" technical idea), or could the RS somehow need some other assurances that it can be compliant? This goes right back to the Adrian clause, and complementarily, to the Shoebox idea. The RS could either warn Alice, or could verify the requesting party on its own. Is this broad-based, or truly healthcare-specific? To how many jurisdictions does it apply? Adrian suggests another case of a warning. A data processing process might take some time to complete, and only licensed practitioners can receive in-process results.

Would it be useful to develop a model clause, or even a pair of model clauses? Sounds like it's useful to try, and then go back and see if it's getting used.

What does the "RO access limit discretion" use case mean? There is something of an opposite to the Adrian clause in the same section: "The resource server MUST NOT give access in the case of an invalid RPT or an RPT associated with insufficient authorization."

A way of looking at it is:

- Conditions: RO wants sharing to happen; AS executes those directions appropriately, such that the RPT is valid and is associated with sufficient authorization data:
    - RS gives access - normal case
    - RS doesn't give access
        - To remain compliant with the law of some jurisdiction - MUST PRODUCE NOTIFICATION?
        - For some other reason - MUST NOT HAPPEN WITHIN AN UMA FLOW GOVERNED BY THE PAT MINTED WITH THE AS ABOVE?
- Conditions: RO doesn't want sharing to happen; AS executes those directions appropriately, such that the RPT is invalid or is associated with insufficient authorization data:
    - RS doesn't give access - normal case
    - RS does give access
        - To remain compliant with the law of some jurisdiction - MUST PRODUCE NOTIFICATION?,
        - For some other reason - MUST NOT HAPPEN WITHIN AN UMA FLOW GOVERNED BY THE PAT MINTED WITH THE AS ABOVE?

The idea of getting notices is that the RO could severe the relationship with the RS if they want. Of course, sometimes there are legal constraints on even getting a notification.

Can you ever trust the AS to serve the RS, versus the RO? Only at the business agreement level could an AS start to consider jurisdictional issues that RS's care about.

Please note: No meeting next week! We'll resume the week after.

2016-01-08

- Review proposed mission and timeline
- Review high-level requirements
- Review latest candidate model text

Attending: Eve, Ann, Jim, Adrian, Jon, Mark, Thomas, Bill Wendell (guest - ears only), Andrew Hughes, Sal

Mission and timeline discussion

A proposed mission and timeline were sent in this message.

Jon speaks in support of the three goals. Is "model clauses" the right phrase, given that "EU model clauses" is something of a term of art? We're free to call it whatever we want. RoboClauses? 🙂 Maybe the connection to EU model clauses is helpful to us.

Regarding meeting scheduling, if text review is paramount in the early going, availability of interested parties is the most key. Offline review and reporting of comments by others could work great.

"Model text" was Eve's temporary phrase covering not just the clauses but also the definition. "Model T – you can start with any color you want, as long as it's black!"

We're thinking that the "standard" meeting time of Fridays at 8am PT would be okay for starters in working on the model text. Note that Eve can't make January 22, so maybe we can move or skip that one.

Andrew's project is holding a F2F in January, and by February will possibly be really ready to need our outputs. Right now they're wrestling with different views on what "authorization" means. Is it authorization to release information? access to an API? (UMA would rely on the semantics of the API being protected for the definition.) Mark notes that consent and permission definitions should come into the picture as well. This does actually impact our term definitions because we use the word "Authorization" as part of our UMA role names, and so we may have to define Authorization, or at least define it in a limited sense so that a context in which the word is defined in a different sense is not incompatible or confusing wrt our usage. If there is, somewhere (EU?), a standard vocabulary, can we leverage it? Or is it too backwards-looking?

Let's add this to our open issues.

AI: Eve: Add model text issues to our issues backlog and tag them appropriately, including the issue of defining authorization, consent, etc.

Consensus on the mission is positive. What should our timeline actually be? Jim recommends that we target a real version number of 1.0. We can achieve this sooner rather than later if we have real review by people who really need this text!

Eve has been collecting external reviewer candidates. We did some more brainstorming.

The right time for a beta version of the model text for review: end of March 2016.

AI: Eve: Send continuing calendar invitations.

Requirements discussion

Some proposed high-level requirements were sent in this message.

Eve did a quick reading of the proposed high-level requirements. People started dropping off, but it sounded like there was at least some weak rough consensus. 🙂 We'll work with this list for now until it proves problematic.

Draft model text

The latest model text for review, with commentary, was sent in this message. It lives persistently here.

We took a quick look at the message sent out.