

UMA telecon 2019-01-10

UMA telecon 2019-01-10

Date and Time

- **Thursdays 6am PT (new time for the new year)**
 - Screenshare and dial-in: <https://global.gotomeeting.com/join/857787301>
 - See UMA calendar for additional details: <http://kantarainitiative.org/confluence/display/uma/Calendar>

Agenda

- Roll call
- Meeting logistics
 - No call on Jan 24 or Jan 31
- Approve minutes of UMA telecons [2018-12-06](#), [2018-12-13](#), [2018-12-20](#)
- Last call: [Identiverse call for presentations](#) is open till Jan 11
- High assurance of the RO to the RqP
 - Is Peter available to describe his method of strong binding from Alice registration time to ensure "high RO IAL" in more detail?
- UMA business model
 - Draft report work
- IETF 104 plans
 - [Resource-indicators](#) draft relevance either way?
 - Next steps?
- AOB

Minutes

Roll call

Quorum was not reached.

Meeting logistics

- No call on Jan 24 or Jan 31

...unless we've got lots of things under discussion and we've got a chair pro tem available.

Approve minutes

- Approve minutes of UMA telecons [2018-12-06](#), [2018-12-13](#), [2018-12-20](#)

Deferred.

Last call: Identiverse call for presentations is open till Jan 11

Eve has one UMA-related submission in the works.

High assurance of the RO to the RqP

- Is Peter available to describe his method of strong binding from Alice registration time to ensure "high RO IAL" in more detail?

Peter isn't available today.

Alec thought about the general CIBA-ish topic and wondered if the AS could use request_submitted error path somehow to leverage CIBA to authenticate the RO and assure the RqP of the RO's identity as required. George notes that there has been discussion in CIBA about a "questioning API", a la the old [Liberty Interaction Service](#). Eve asks: Isn't identity assurance of some sort the entire point of OIDC on top of the authorization endpoint? George: Except that this is why CIBA is about a "back-channel" flow. Hmm, note that what they are calling "back-channel" is more like "out-of-band", without the use of redirection. They don't mean asynchronous. There is a "front-channel" interaction with a device that the user in question was never redirected to. It's an attempt to synchronize using an asynchronous method. The benefit of UMA is supposed to be that it's prepared to go totally asynchronous. The request_submitted error enables some kind of synchronous check with the RO (which perhaps we could leverage for a forceAuthn check of Alice if we can figure this out).

Can CIBA be turned into a clean Interaction Service that we could then use for the specific places where it would be valuable, such as the AS checking the RO and the AS checking the RqP? What questions can be asked of the user (RO in our use case)? We discussed the concept of a "binding message" (see [binding_message](#) defined [here](#)), but its use seems to be limited more to a type of one-time code that defeats an MITM attack or similar that disrupts messaging between the consumption and authentication device.

Are there privacy considerations in the RqP learning the RO's identity information or authentication level at any point before the RqP has been granted access to the resource? If the RqP has a condition such as "the RO must have authenticated to level X in the last N minutes", then presumably those conditions along with the RO's policy conditions must all obtain before access should be given. Might this affect the UMA set math as well? Potentially, yes.

AI: Eve: Ask Bjorn about availability of GSMA Mobile Connect use cases that get into more detail than the [CIBA section](#). Is [this link](#) in the right direction? George thinks there's a document that is 30-40 pages in length.

Nancy notes that the Apple Health app seems to ask for access to medical records once, but then gets updates continually thereafter. It's likely OAuth-based, where she can withdraw consent for that app connection, pairwise. The model still seems to be "Alice-to-Alice sharing". If Alice were to want to share a record with Dr. Jones, we now get into use cases like: Alice wants Dr. Jones to prove he's himself once a month (classic UMA); Dr. Jones wants Alice to prove that the records from a specialist that she's sharing with him are truly hers to share (the new CIBA-ish use cases we're contemplating)...

UMA business model

- Draft report work

Eve got some work done on her draft, with help from Domenico on new diagrams. She and Tim will try and get together before our next call and press ahead.

IETF 104 plans

- [Resource-indicators](#) draft relevance either way?
- Next steps?

The resource-indicators draft seems to mean both "resource server" and "protected resource". UMA allows for registration of resources, which spells out the specifics. George suggests that this spec is trying to split the difference in management overhead level for high-value resources and low-value resources in splitting the difference in the meaning. We had previously discussed (but, at the time, discarded) the idea of having "wildcards" (need to find the specific GitHub issue) when registering resources, so that you don't have to fully qualify resource names when registering them. Lots of types of services might have millions of resources – or put another way, an infinite number of dynamic resources, depending on how they're created, torn down, etc. Having to register them in a fully qualified way may be too static, too heavyweight, etc.

The spec binds access tokens to a particular URI. Because most aren't downscoping their tokens, leakage of tokens will enable too-great service access for its lifetime. This binding is like an audience check, only for the original service it was meant for. UMA's equivalent, in part, is the nature of the RPT and set math, and the fact that it [represents](#) an intersection of who wants it and who's granting it. The nature of policy conditions makes the RPT a bit more "sender-constrained" (that is, constrained by what the RO wants to happen) by definition, if not in construction. However, the new reverse-proxy phishing attacks are basically stealing session cookies, which the RPT is a bit like. We should take a look at any new mitigations we should be applying along the lines of true "sender-constraining".

Discussion for next time:

- Sender-constrained RPTs?
- Comments on OAuth-distributed draft?
- Contributing the UMA2 specs?
- Attending IETF 104?

Attendees

As of 18 Oct 2018, [quorum](#) is 5 of 8. (Domenico, Peter, Sal, Andi, Maciej, Eve, Mike, Cigdem)

1. Domenico
2. Sal
3. Eve

Non-voting participants:

- Alec
- George
- Nancy
- Tim

Regrets:

- Peter
- Cigdem
- Mike