OAuth 2.0 Token Exchange

An STS for the REST of Us

Brian Campbell
et al.

IETF 95
Buenos Aires
April 2016

https://tools.ietf.org/html/draft-ietf-oauth-token-exchange-04

Circa Early 2015

- draft-ietf-oauth-token-exchange-02
  - By Mike Jones & Tony Nadalin
- draft-campbell-oauth-sts-02
  - By Brian Campbell & John Bradley

‘Discussed’ in Prague…

(no heads were severed)

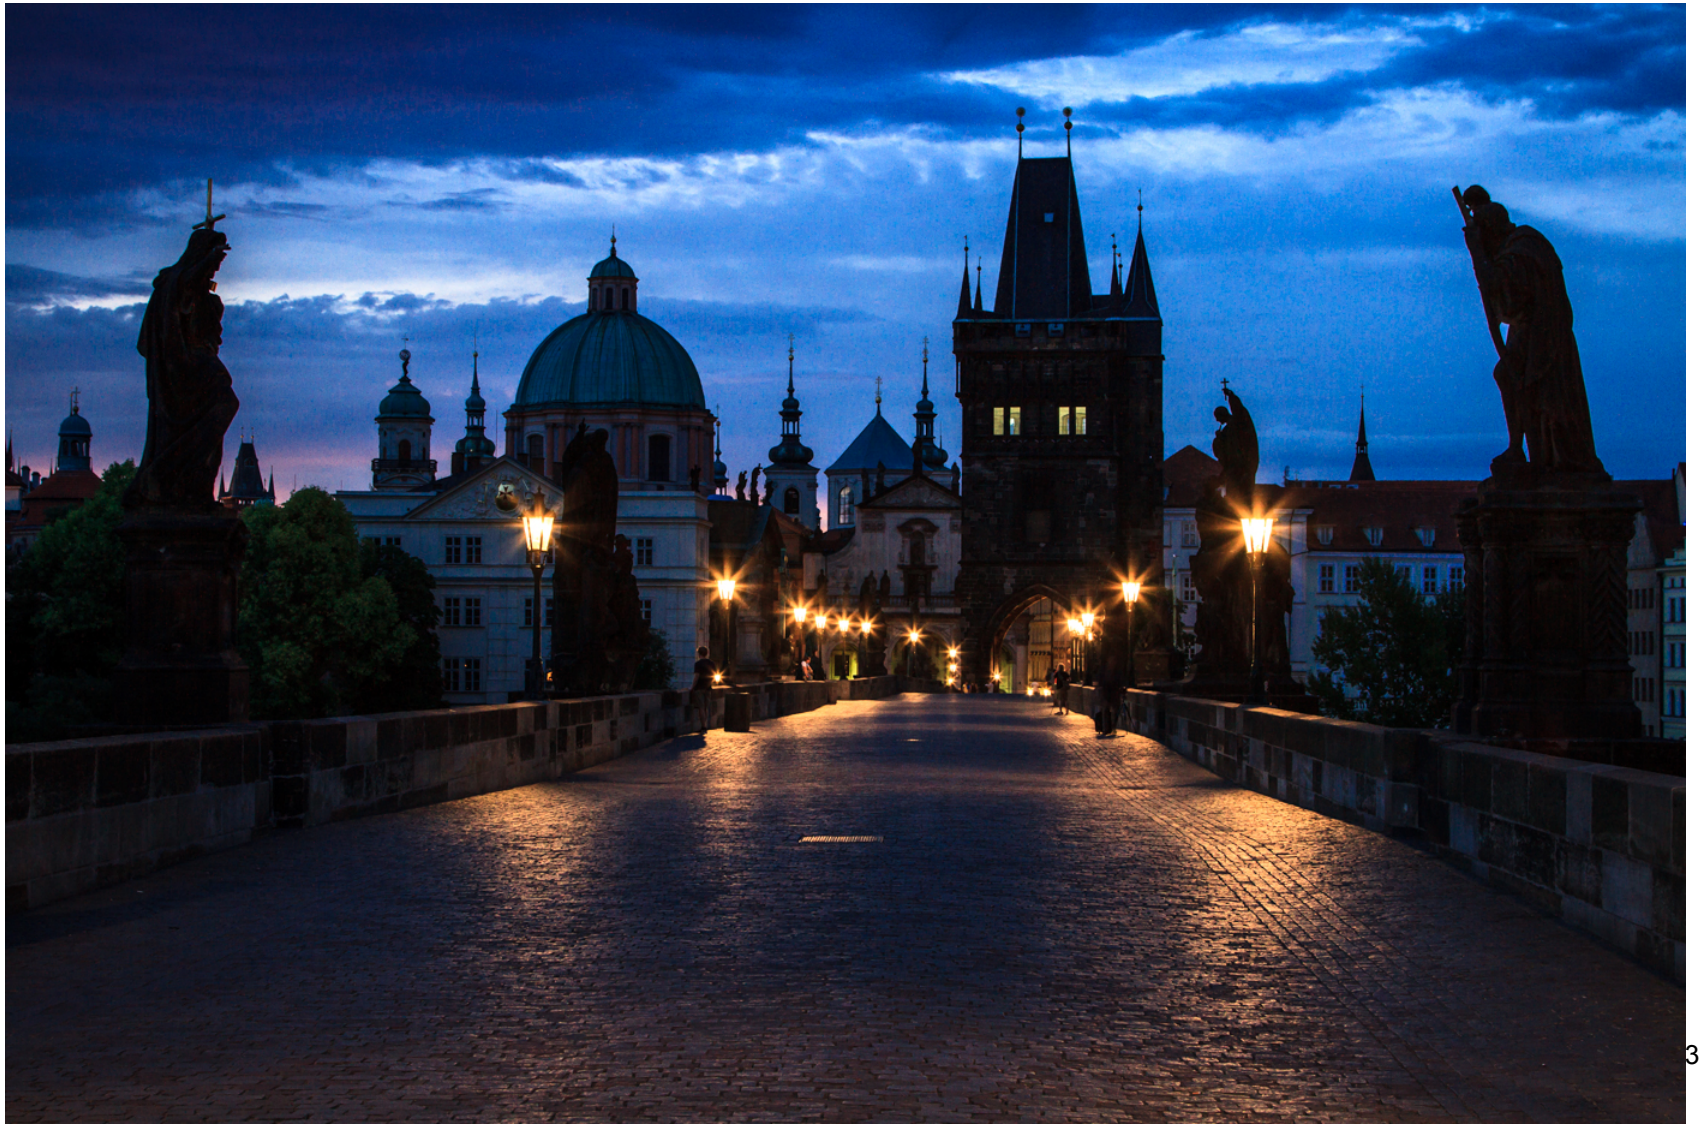

Reconciliation

- http://tools.ietf.org/html/draft-ietf-oauth-token-exchange-03 in Dec ‘15
  - Reconciled the differences the two drafts & incorporated WG rough consensuses input
  - Changes listed at: http://tools.ietf.org/html/draft-ietf-oauth-token-exchange-03#appendix-D
- https://tools.ietf.org/html/draft-ietf-oauth-token-exchange-04 in March ‘16
  - Clarifications in text and examples
  - Defined and registered urn:ietf:params:oauth:token-type:id_token
- Too many changes to discuss in time allowed
  - Please read the draft, if you haven’t yet

Open Issues

- Should there be a way to use short names for some common token type identifiers? URIs are necessary in the general case for extensibility and vendor/deployment specific types. But short names like access_token and jwt are aesthetically appealing and slightly more efficient in terms of bytes on the wire and url-encoding. There seemed to be rough consensus in Prague ('No objection to use the proposed mechanism for a default prefix' from https://www.ietf.org/proceedings/93/minutes/minutes-93-oauth) for supporting a shorthand for commonly used types - i.e. when the value does not contain a ":" character, the value would be treated as though urn:ietf:params:oauth:token-type: were prepended to it. So, for example, the value jwt for requested_token_type would be semantically equivalent to urn:ietf:params:oauth:token-type:jwt and the value access_token would be equivalent to urn:ietf:params:oauth:token-type:access_token. However, it was a fairly brief discussion in Prague and it has since been suggested that making participants handle both syntaxes will unnecessarily complicate the supporting code.
- Provide a way to include supplementary claims or information in the request that would/could potentially be included in the issued token. There are real use cases for this but we would need to work through what it would look like.
- Understand and define exactly how the presentation of PoP/non-bearer tokens works. Of course, the specifications defining these kinds of tokens need to do so first before there is much we can do in this specification in this regard.
- It seems there may be cases in which it would be desirable for the authenticated client to be somehow represented in the issued token, sometimes in addition to the actor, which can already be represented using the act claim. Perhaps with a client_id claim?
- (from the list) Should the act and may_act also be registered for Introspection Endpoint responses?