OSCORE Profile of ACE

https://tools.ietf.org/html/draft-ietf-ace-oscore-profile-08

Francesca Palombini, Ericsson AB
Ludwig Seitz, RISE SICS AB
Göran Selander, Ericsson
Martin Gunnarsson, RISE SICS AB

Status

- Waiting on Ben’s go-ahead to submit new version:
  - PR: https://github.com/ace-wg/ace-oscore-profile/pull/25
  - Issues: https://github.com/ace-wg/ace-oscore-profile/issues/26
- Diff with last submission:
  - https://tools.ietf.org/rfcdiff?url1=https://tools.ietf.org/id/draft-ietf-ace-oscore-profile.txt&url2=https://ace-wg.github.io/ace-oscore-profile/v-09/draft-ietf-ace-oscore-profile.txt
- Some open points discussed here and posted to the mailing list:
  - https://mailarchive.ietf.org/arch/msg/ace/Edjf_Pb0rDK1LoZChBE03ZQ9ajE

Open points

- “I know it when I see it” – HMAC based HKDF and AEAD COSE algorithms
  - 43, 44
- Remove replay window from the OSCORE_Security_Context object
  - *The idea here was that the AS is able to tell the RS what type and size of replay window to use, but I do not see a reason why the AS should send this info to the RS. Window type and size is implementation related, and each RS node can have its own.*
  - 46
  - Done in PR
- Profile indicates that C/AS and/or RS/AS use OSCORE, or that information is not expected to be encoded in any protocol element?
  - 92
  - (Next slide)

Open points - 92

92.   o Profile Description: Profile for using OSCORE to secure
> >   communication between constrained nodes using the Authentication
> >   and Authorization for Constrained Environments framework.
> >
> >   This is the registry for ACE profiles; "using the Authentication and
> >   Authorization for Constrained Envirornments framework" feels redundant.
> >   There's also perhaps something to say about defining "coap_oscore" for all
> >   three interaction flows even though the procedures for two of the three are
> >   essentially just "out of the scope of this specification", though having
> >   *some* identifier for "use OSCORE" is probably still useful.
> >
> > FP: Ok, we can remove the last part. About the interaction flows between nodes, I am not sure what you mean, as by following definition https://tools.ietf.org/html/draft-ietf-ace-oauth-authz-31#section-5.6.4.3 it is not mandatory for a profile to specify C-AS and RS-AS, those flows can be out of scope. Am I missing something?
>
>   It's not mandatory, though given that we talk about profiles not being able
>   to assume that the same profile is used for any/all of C/RS, C/AS, and
>   RS/AS, I guess I assumed that we use the profile identifier to talk about
>   the last two cases as well as the first one. Maybe I'm wrong to assume
>   that!
>
> FP: From my understanding, that's not the case, as I said above. Maybe someone else can confirm.
>
>   Regardless, I think the question here is whether we want to have people use
>   the "coap_oscore" string to indicate that C/AS and/or RS/AS use OSCORE, or
>   if that information is not expected to be encoded in any protocol element;
>   if the latter, then there's nothing to change.
>
> FP: As of now, that information is not expected to be encoded in any protocol element, as far as I understand. Again, if anybody can confirm/shout if wrong, please do.

Okay. I would be happy to hear from the WG if there are other known uses
like this.

Open points

- The mechanism of letting the RS pick the identifier of the client is not worth the additional complexity.
    - *"[JLS] For better or worse, this is not the only consideration. The above statement presumes that only the AS is going to be assigning the client identifiers. If the RS uses a different method of doing authentication as well, such as LAKE, then there is a chance for collisions at that point as well."*
    - Define an error response RS-C if Id collision?
    - 6, 7, 32, 61, 65

- Recommendation about length of nonces N1 and N2 to use.
    - *" [JLS] I am agnostic on this, but I think that generally 64-bits from each side is sufficient. This does add an additional 8 bytes of message size when sending the token to the server. That worries me slightly.*
    - *"I agree that 64 bits from each endpoint may well be sufficient, but it requires some analysis and thought to verify that, whereas 128 bits is "obviously enough". So, I'd like to know what level of analysis has already been done (vs. just inference from, e.g., OSCORE)." - Ben*
    - 5, 52

- Define and register 2 new ACE parameters to transport the nonces used in the exchange, instead of using "cnonce".
    - 3, 53, 60