TEFCA Facilitated FHIR Implementation Guide
Draft 1
Published: October 7, 2022
Feedback from the eHealth Exchange
Prepared by Eric Heflin, Consultant

Introduction

We appreciate the opportunity to provide feedback on the draft version of this new Implementation Guide. Overall, we feel the IG is quite clear, complete, well-written and serves to advance the industry. After careful review, we respectfully offer the following suggestions for improvement which we sincerely hope assists the RCE in its efforts. The eHealth Exchange has been using FHIR in validation and production for the FDA for over a year and we are freely sharing our relevant lessons learned in this document with the hopes of helping the country.

Any errors are the responsibility of the author, not the eHealth Exchange.

Overall Comments

We respectfully request that the IG includes conformance statements for all testable requirements, like the various "QTF-###" conformance statements in the QHIN Technical Framework. These conformance statements ideally should be directly translated to test statements, and, in totality, an entity that passes all these conformance statements should be interoperable and conformant to this specification.

Please consider adding push FHIR transactions as well as query/response transactions.

Are FHIR retrieve operations in scope? If so, can that be clarified via a statement indicating such.

Are FHIR intermediaries prohibited or allowed? The document is silent on this consideration. The eHealth Exchange urges the RCE to specifically state that FHIR intermediaries, both passive and active, are optional and are permitted by this IG. FHIR intermediaries have the potential for providing value to the FHIR initiator and responder actors. For example, the eHealth Exchange Hub is used by the FDA as an active component to mediate the differences between various data holder FHIR server behaviors for the FDA COVID-19 Adverse Events Case Follow Up production use case.

It may be impractical for a QHIN to also be a Certification Authority as noted in more detail below.

We respectfully request that the IG clarify which version, or versions, of FHIR are supported, and the associated policy for subsequent versions of FHIR (are subsequent FHIR versions optional, prohibited, mandated after a phase-in period, etc.).

We applaud the use of dynamic registration since it is highly valuable and should remain mandatory for all TEFCA Participants to enable automated scalability.

Our experience in using FHIR in production is that all current responding EMR vendors hard-code the scopes and purpose of use at the time the client_id is statically registered. We appreciate how this IG mandates that both OAuth 2 scopes and purpose of use can be determined dynamically at run-time.

OAuth 2.0 has been superseded by the OAuth 2.1. Version 2.1 is a non-breaking update that remediates several problems with OAuth 2.0. The RCE should consider adopting OAuth 2.1, especially at the present time where this is a "green field" and there are no existing OAuth 2.0 implementations; It will be harder to adopt OAuth 2.1 in the future.

In all cases where responses to FHIR requests, or OAuth requests, are not specified already, we respectfully request that the RCE constrain the error/exception responses to a specific value-set that includes both human-readable and programmatically processable fields. The RCE should curate this list of exception handling values and extent in in the future via a public process. If the RCE does not constrain exception handing in this way, then each QHIN or data holder is likely to implement exception handling differently, harming interoperability.

Regarding Section 5.2 Authentication/Trust

Certification Authority (CA)

The risks of requiring a QHIN to manage their own CA is extremely high and, the costs are high, and if the risks are realized, could undermine trust in the TEFCA due to a misstep by any of the QHIN CAs.

While admirable that the RCE envisions that QHINs can be empowered to issue their own End Entity X.509 certificates, its essential that QHIN CA functions be held to industry standards. One (of many) example documents that prescribes a **fraction** of these issues can be found in the industry CA/Browser Forum Baseline Requirements at: https://cabforum.org/wp-content/uploads/CA-Browser-Forum-BR-1.8.4.pdf which is based on an industry-standard template.

Operating a CA is a very demanding endeavor, requiring hardened infrastructure, plus hundreds of other requirements including:

- Vetted staff with experience establishing and operating a CA
- Hardware Security Modules
- Legal review
- Operational security
- Policy and procedures
- Transparent governance
- Web applications to administer the system
- Web applications for end-user certificate acquisition and management
- Security specific X.500 LDAP directory
- Scalable CRL list publication
- Scalable OCSP responder network
- And more

An alternative approach would be for the RCE to create a list of pre-approved qualified organizations that a QHIN can select from, to manage that QHIN's CA functionality.

For these reasons, we respectfully suggest that the RCE mandate that QHINs use a vetted, pre-approved, CA instead of the QHIN acting as the CA themselves.

Regarding: 5.2.1.2 Issuance

If one Operator maintains a client or server for multiple organizations, how will the partners of this client or server know who their exchange partners are? Will some aspect of the exchange, or the X.509 certificates, discern among the various organizations? Example: Let's say a given OperatorX maintains a client for organizations A, B and C and used the same X.509 certificate for all three organizations. If the OperatorX client initiates a FHIR outbound request, how will the responder know if the request came from A, B, or C? If this discernment is to be done via HL7 UDAP JWT attributes, then can the 5.2.1.2 section be updated to reference that HL7 UDAP attributes MUST be different, and specific, for requests from organization A than organization B or C even though the X.509 certificate is the same for all three organization?

Regarding: 5.2.1.3 Structure

What is the behavior when and if more than one certificate is issued with the same Application name? For example, what if the name is re-used for different versions of that Application over time? We recommend that the RCE require the Common Name to be unique across all X.509 certificates issued by that CA. (I believe most CA systems already require such.)

Regarding: Software Statement

The text "The client signs the software statement using one of the RS256, ES256, RS384, or ES384 signature algorithms as defined in RFC 7518; the algorithm used will depend on whether the client app's X.509 certificate contains an RSA or EC key. All implementations SHALL support RS256, SHOULD support ES256, ES384 and RS384" has an apparent logic error: In order for the Client X.509 certificate to support RS256, as a mandatory baseline, then the certificate in all cases must contain RS256 keys. So, the use of ES256 or ES384 would dictate that the certificate contains BOTH RSA and EC keys.

We respectfully suggest that the text be re-worded to something similar to: "The client determines which of the signature algorithms are supported in common by the client and the server, and then signs the software statement using one of the RS256, ES256, RS384, or ES384 signature algorithms as defined in RFC 7518; the client SHOULD chose the strongest mutually supported signature algorithm. All implementations SHALL support RS256, and SHOULD support ES256, ES384 and RS384. All X.509 certificates SHALL have RS256 keys and SHOULD have ES256, ES384 and RS384 keys".

Regarding: Inclusion of Certifications and Endorsements

The RCE defines the return of HTTP response code of 201 as indicating a successful registration. The RCE should also define the HTTP response codes for exceptions and errors to ensure interoperability of these responses. As mentioned above, these codes should be curated by the RCE, and should have mandatory human-readable and mandatory programmatically processable fields.

Regarding: 5.2.3.2 TEFCA Basic App Certification Profile

The draft text states "Authorization servers that reject a registration request due to a missing element SHOULD respond with an informative error identifying that element." The RCE should not leave the list of possible exceptions up to each Authorization Server, because the result will be that each Authorization Server will implement different response codes and logic, and thus this will not be interoperable. Instead, the RCE should include a list of mandatory responses for specific error conditions and then the RCE should curate that list of responses to maintain interoperability over time. Furthermore, the RCE should mandate a computable vocabulary for the exception response structure with a required computable, and a required human-readable, value. Finally, for the success case (HTTP 201 response), the RCE should precisely define the returned client_id value in the HTTP 201 returned response (where is the Client ID returned).

Regarding: 5.2.3.3 Modifying Registrations

Regarding the text "If the Authorization Server returns a different client_id in the registration response, the client application SHALL use only the new client_id in subsequent transactions with the Authorization Server." The RCE should mandate that the responding Authorization Server disable the old client_id so that it cannot be used for subsequent requests (for increased security). However, the retired client_id should be preserved by the Authorization Server so that it can be associated with log entries and the requester.

Regarding: 5.2.4 Authorization Code Grant Type (3-legged OAuth 2.0)

UDAP.org has no standing. The RCE should not reference UDAP.org but should reference true ISO/ANSI standards bodies and standards such as those published and curated under IHE or HL7. Pointing to a private web site, such as UDAP.org, makes this document vulnerable to the preferences of that private organization, and offer no assurances provided by industry standards bodies such as open, transparent, neutral standards development and maintenance over time conducted under oversight ensuring anti-competitive regulations are observed. UPDAP.org provides no such assurances.

Regarding: 5.2.4.2 Obtaining an Access Token

The draft text reads "If the organization operating the requesting application has additionally identity-proofed the end user of its application, then the requesting application MAY provide metadata about the user to the data holder as additional authorization information …" We respectfully suggest that the RCE should change this from a "MAY" to a "MUST". This will improve accountability and allow for more detailed Accounting of Disclosures responses. In a similar manner, the "tefca_user" key should be mandatory.

Regarding the draft text "The user metadata submitted by the requesting application in the extension object SHALL correspond to the verified identity attributes of the permitted user … where the purpose of use code is "REQUEST"), this user is not necessarily the patient who is the transaction subject, i.e., the verified user MAY instead be a patient's authorized representative." There are several high value use cases that the RCE should make mandatorily supported such as release of information for the National Institutes of Health "All of Us" program, where the request is being made by a third party (such as the NIH), but the request is specifically pursuant to a patient authorization. The RCE should mandate that conformant systems support this class of use cases within the scope of the TEFCA. If the intent of the draft text is to enable this type of use case, then the draft text could perhaps be revised such as to say "The user metadata submitted by the requesting application in the extension object SHALL correspond to the verified identity attributes of the permitted user (verified as per Section 3.2) who is making the request. Note that for patient requests (i.e., where the purpose of use code is "REQUEST") where this user is not necessarily the patient who is the transaction subject, i.e., the verified user MAY instead be a patient's authorized representative, or the verified user MAY be a person or a system at organization acting on the specific written authorization of the patient."

Regarding: Table 3 TEFCA User Authorization Extension object

To allow for interoperability, the "lal_vetted" attribute should make the comma-separated list of values a list of prescribe names and values, such as "given_name=john, family_name=doe". Also, some names have comma characters in the same, so there should be a method of escaping the comma such as specifying a backslash character before the command when the comma is part of the patient's name.

Regarding: 5.2.5 Client Credentials Grant Type (2-legged OAuth 2.0)

The draft text states "In this flow, the authorization endpoint is not used". My personal understanding is that the workflow does indeed use the Authorization Server, as per the diagram below. Also, section 5.2.2.2 Client credentials grant at https://build.fhir.org/ig/HL7/fhir-udap-security-ig/b2b.html states (emphasis mine) "Client applications using the client credentials grant and authenticating with a private key and Authentication Token as per Section 5.2.1 SHALL submit a POST request to the Authorization Server's token endpoint containing the following parameters as per Section 5.2 of UDAP JWT-Based Client Authentication."

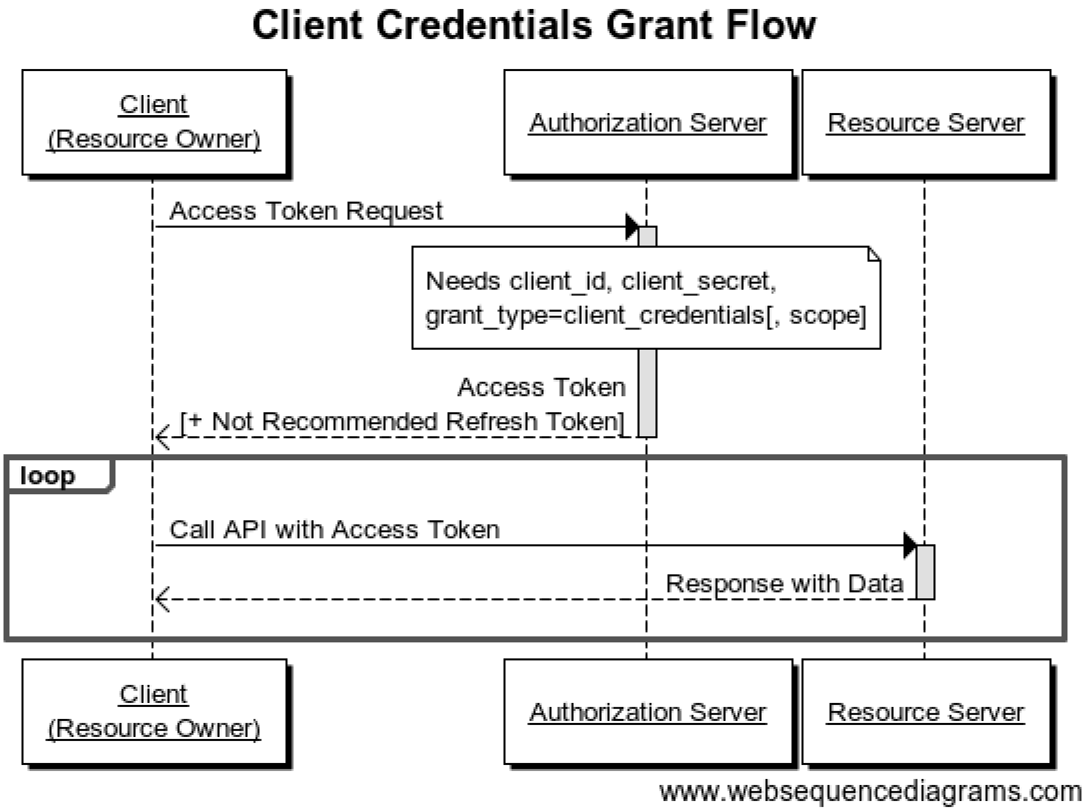
Client Credentials Grant Flow
Client
(Resource Owner)
Authorization Server
Resource Server
Access Token Request
Needs client_id, client_secret,
grant_type=client_credentials[, scope]
Access Token
[+ Not Recommended Refresh Token]
loop
Call API with Access Token
Response with Data
Client
(Resource Owner)
Authorization Server
Resource Server
www.websequencediagrams.com

Can this be clarified?

Regarding: 5.2.6 Individual Access Services (IAS) Requests

Like our comments above, we respectfully suggest that the “REQUEST” purpose of use be expanded to include organizations, such as the National Institutes of Health, making requests on the behalf of a person under that patient’s specific written authorization.

Regarding: 5.3.2 Patient Discovery

The draft text reads “Each query SHALL include, but is not limited to, all available USCDI patient demographics …normalized as per the Project US@ Technical Specification…” Project US@ makes almost all rules optional. We respectfully request that the RCE constrain the Project US@ Technical Specification to indicate additional constraints to make Project US@ implementable.

Conclusion

We appreciate the chance to provide feedback on this draft IG. We sincerely hope that our suggestions improve the next iteration of this important specification.