TEFCA RECOGNIZED COORDINATING ENTITY

Standard Operating Procedure (SOP): Facilitated FHIR Implementation

Version 2.0

Effective Date: March 8, 2026

Applicability: QHINs, Participants, Subparticipants that engage in TEFCA Exchange leveraging Facilitated FHIR

TABLE OF CONTENTS

1 Common Agreement References 3

2 Definitions 3

3 Purpose 3

4 Facilitated FHIR Query Scenario 4

4.1. Actors 5

4.2. Assumptions 5

4.3. Pre-conditions 6

5 Use Case Steps 7

6 Procedure 9

6.1. Overarching Requirements 9

6.2. FHIR Security 10

6.3. Exchange Partners 10

6.4. Exchange Purposes 11

6.5. General Requirements 11

6.6. FHIR Endpoints & Endpoint Discovery 11

6.7. Patient Matching 12

6.8. Provenance Use 13

6.9. Error Responses 13

6.10. Security 13

6.11. Requirements for Use of HL7 SSRAA 14

6.12. Scope Negotiation Requirements 17

7 Version History 19

1 Common Agreement References

The requirements set forth in this Standard Operating Procedure (SOP) are for implementation, in addition to the terms and conditions found in the Framework Agreements, the Qualified Health Information Network™ (QHIN™) Technical Framework (QTF), and applicable SOPs. The Trusted Exchange Framework and Common Agreement™ (TEFCA™) Cross Reference Resource identifies which SOPs provide additional detail on specific references from the Common Agreement.

All documents cited in this SOP can be found on the Recognized Coordinating Entity® (RCE®) website.

2 Definitions

Select terms used throughout this SOP are defined in this Section for ease of reference. All capitalized terms used in this SOP have the respective meanings assigned to such term in the TEFCA Glossary.

FHIR Adopters: Any QHIN, Participant, or Subparticipant that wishes to engage in TEFCA Exchange leveraging Facilitated FHIR, as described in this SOP.

3 Purpose

This SOP identifies specific requirements for Facilitated FHIR implementation activities. Any FHIR Adopter may participate in any of the FHIR activities specified in this SOP.

The goal of this SOP is to encourage consistent adoption of a scalable, network approach to Facilitated FHIR across TEFCA. This SOP provides a roadmap that allows for a transition period to ease adoption of an eventual network-wide approach. Many of the conformance requirements in this SOP are optional to allow for maximum flexibility during the transition period. This will likely lead to the need for direct coordination between FHIR Adopters and could result in acceptable inconsistencies in the scope of data exchanged, the Exchange Purposes selected for Facilitated FHIR implementation activities, and the initial partners with whom each FHIR Adopter exchanges. The FHIR Implementation Advisory Group[1] is established to collect and document learnings and progress towards this goal, including assessing adoption and implementation of the HL7 Security for Scalable Registration, Authentication, and Authorization (HL7 SSRAA) FHIR Implementation Guide (IG). The FHIR Advisory Group will provide regular updates to the TEFCA Transitional Council or the Governing Council (as applicable) on progress, and informed by that progress, will have the authority to recommend changes to the dates in the FHIR Security Roadmap detailed in Section 6.2 of this SOP and/or conformance statements throughout this SOP for approval by ONC and the RCE.

[1] See Advisory Groups SOP: https://rce.sequoiaproject.org/tefca-and-rce-resources/

4 FACILITATED FHIR QUERY SCENARIO

In this scenario, a health care provider treats a patient in an emergency department and seeks to retrieve information regarding the patient's care from the patient's primary care provider(s) through Facilitated FHIR TEFCA Exchange.

The basic pattern of the flow follows the Patient Discovery and Document Query flows and then diverges to use FHIR queries to identify the specific patient and query for specific FHIR resources.

Once the Initiating Node has the appropriate endpoints it begins a HL7 *FAST* UDAP Security for Scalable Registration, Authentication, and Authorization (SSRAA) Trusted Client Registration to assert its identity to the authorization server using a TEFCA certificate. Once identified and issued a client_id, the Initiating Node authenticates, authorizes access, and receives an access token.

Once authorization has been granted, the Initiating Node queries the FHIR server for the appropriate Patient Resource including the demographics known to the Initiating Node and begins to Query for that patient's health care data.

Note: This flow assumes SSRAA IG use.

Specified standards for a Facilitated FHIR Query are included in **TABLE 1. SPECIFIED STANDARDS FOR FACILITATED FHIR QUERY**.

TABLE 1. SPECIFIED STANDARDS FOR FACILITATED FHIR QUERY

Query Functions	Specified Standard(s) / Profile(S)
Time Management	• IHE Consistent Time (CT)
Secure Channel	• IETF TLS 1.2 w/ BCP-195[2] or • IETF TLS 1.3 w/ BCP-195
Node Registration	• OAuth V2.0 • HL7 *FAST UDAP* Security for Scalable Registration, Authentication, and Authorization V1.1.0[3]
User Authentication	• OAuth V2.0 • HL7 *FAST UDAP* Security for Scalable Registration, Authentication, and Authorization V1.1.0
Authorization & Exchange Purpose	• OAuth V2.0 • HL7 *FAST UDAP* Security for Scalable Registration, Authentication, and Authorization V1.1.0 • SMART Application Launch Framework Implementation Guide Release 1.0.0: **SMART App Launch: Scopes and Launch Context** or higher versions

[2] *Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)* (IETF BCP 195) available at https://tools.ietf.org/html/bcp195.
[3] https://hl7.org/fhir/us/udap-security/STU1.1/

Query Functions	Specified Standard(s) / Profile(S)
Query for Patients	• IHE XCPD • FHIR R4 V4.0.1 • HL7 FHIR US Core Implementation Guide V6.1.0 or higher
Information Query and Retrieve	• FHIR R4 V4.0.1 • HL7 FHIR US Core Implementation Guide V6.1.0 or higher • FHIR Implementation Guides as required by the Exchange Purposes (XP) SOP and XP Implementation SOPs, as applicable
Auditing	• IHE ATNA (QHINs; Content only) • ASTM E2147-18 (Participant/Subparticipant; Content only)

4.1. Actors

The following lists the Actors and services included as part of the workflow. Cardinality represents the number of that Actor/service expected and which QTF "system" Actor is expected to have that service or Actor role. The Initiating Node uses a FHIR $match operation with the Patient Discovery demographics to validate the Response from the QHIN and to gain patient context.

Actors/Services	Cardinality	System Actor
Initiating Node	1..1	Any Initiating Node
Initiating Gateway	1..1	Initiating QHIN
QHIN Directory	1..1	Initiating QHIN
QHIN Directory	1..*	Responding QHIN(s)
Responding Gateway	1..*	Responding QHIN(s)
Responding Node(s)	1..*	Any Responding Node

4.2. Assumptions

1. All Initiating and Responding Nodes agree on transport level details (specified for transactions between QHINs elsewhere in this document) that allow for the following:
   a. System authentication and encrypted communications over a secure channel;
   b. The ability to provide information in each transaction that identifies security and permission details about the Query such as: who is sending, what their role is, and what their Exchange Purpose is; and
   c. The ability of Actors to choose if/how to allow a transaction to proceed based on privacy policies, security details, and the requirements of the Common Agreement.
2. The Initiating Node may not know the patient identifier(s) and/or Responding Node(s) for a Query.
3. The RCE Directory has up-to-date listings of all FHIR-capable Participants' and Subparticipants' FHIR Endpoints.

4.3. Pre-conditions

This workflow assumes the following conditions:

1. The Initiating Node knows sufficient patient demographics for a successful match as determined by the Responding Node.
2. Each Actor has the appropriate service endpoint(s) and other connectivity information for any other Actors above or below it in the hierarchy with which it connects directly.
3. The RCE Directory includes the organization facility name(s) for all current Participants and Subparticipants.
4. Each QHIN maintains an up-to-date copy of the RCE Directory.
5. Each QHIN has either a Record Locator Service (RLS) or Enterprise Master Patient Index (eMPI) or uses other techniques to perform patient lookup within the Service Level Requirements timeout limitation as specified in the QHIN Service Level Requirements Policy.

5 USE CASE STEPS

Nominal Flow

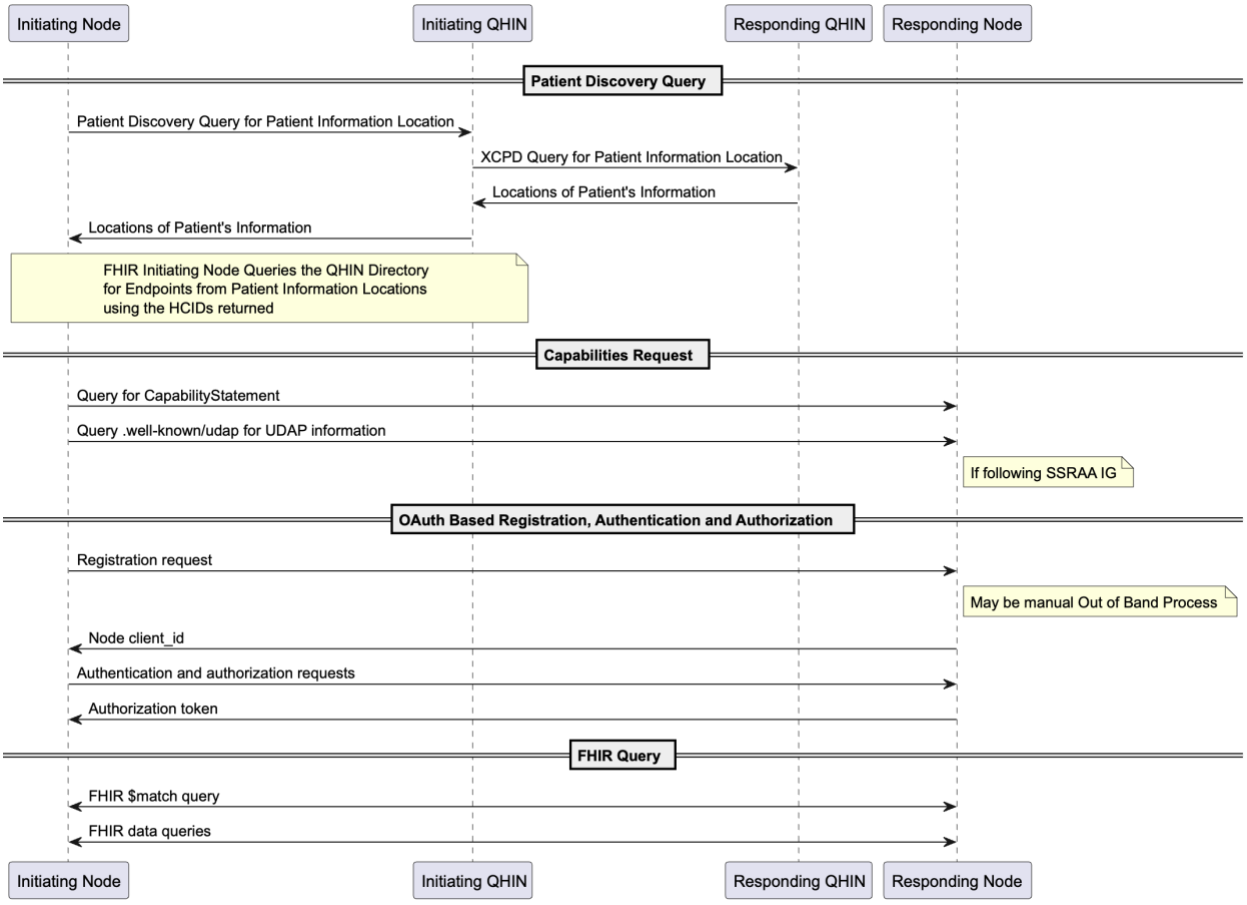

1. The Initiating Node sends a Patient Discovery Query Solicitation, through any intermediary Subparticipants or Participant, as applicable, to the Initiating QHIN to Query for patient information as per the QTF.
   a. The Patient Discovery returns a list of HCIDs with relevant patient information.
2. Initiating node queries the QHIN Directory for FHIR Endpoints using the HCIDs returned.
3. Initiating Node selects which FHIR Endpoint(s) they will be querying for data.
4. The Initiating Node queries for the Responding Node's CapabilityStatement and reviews all capabilities for matches with the Initiating Node requirements.
5. The Initiating Node queries the .well-known/udap endpoint to get all needed information for UDAP registration including the registration endpoint for the Responding Node and supported scopes.

a. If not following the SSRAA FHIR IG, registration is completed out of band, the flow continues at Step 6.

b. The Initiating Node sends a Dynamic Client Registration Request to the Responding Node's authorization server with all relevant information and a list of needed scopes.

c. The Responding Node's authorization server returns a client_id unique to that Initiating Node.

6. The Initiating Node uses the returned client_id, relevant scopes and user information contained in the relevant extensions to Request authentication and authorization to query patient data.

   a. The authorization server grants the Initiating Node a token allowing for querying of data from the Responding Node.

7. The Initiating Node uses the token in the query flow to identify itself to the Responding Node and uses the $match operation with a US Core Patient Resource to gain a list of patients matching the demographics.

8. The Initiating Node selects the appropriate patient from the list provided and begins Querying for associated data.

   a. The Initiating Node and Responding Node create an audit log of all transactions.

Alternate Flow 1: FHIR Push

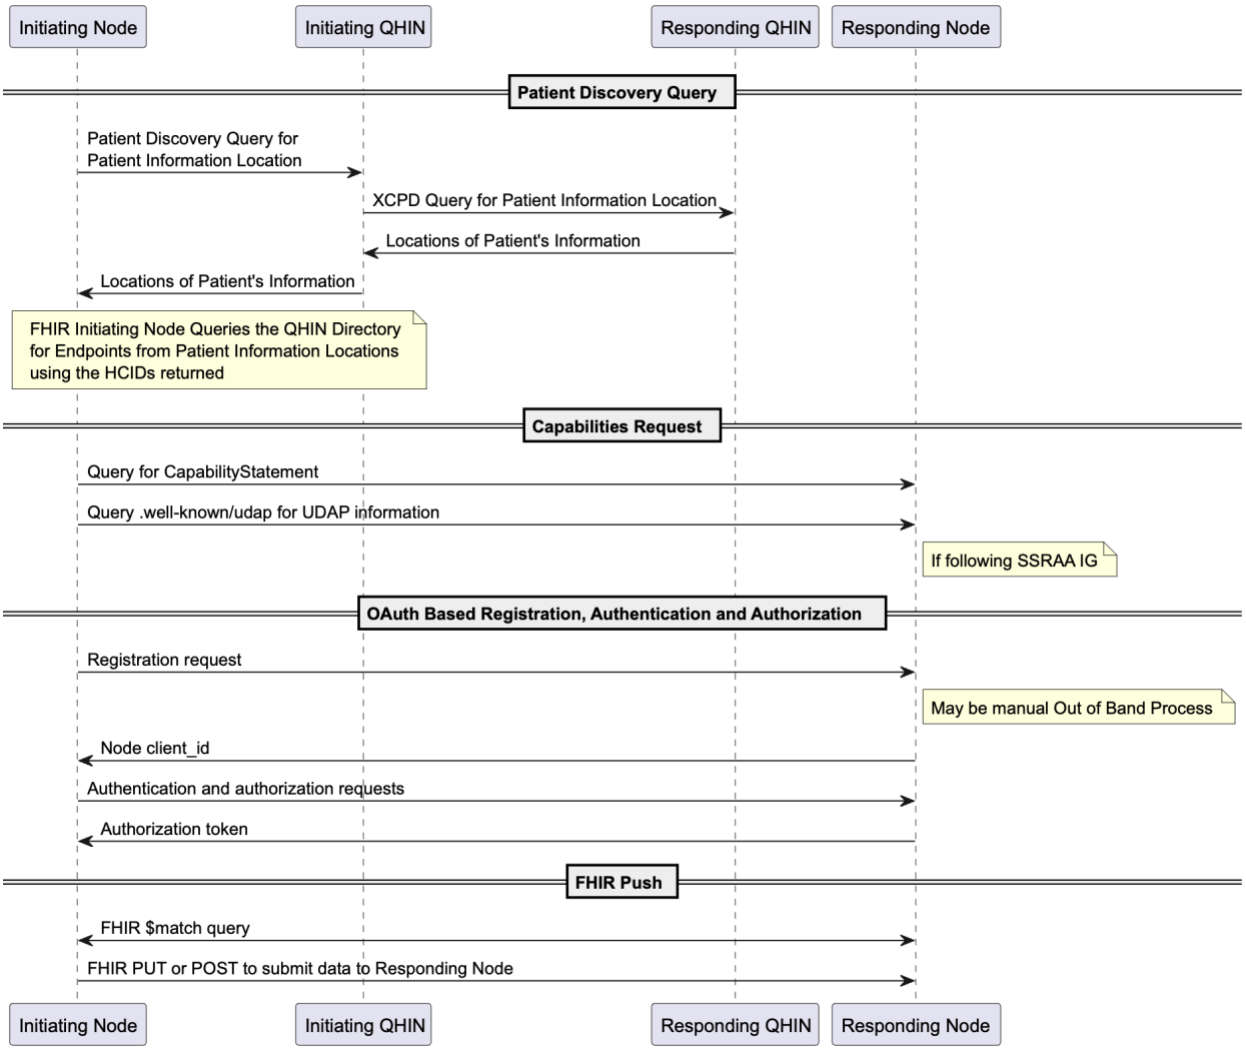

1. Flow begins at step 8 of the Nominal Flow above.
2. The Initiating Node selects the appropriate patient from the list provided and executes a PUT or POST to submit data to the Responding Node.

6 PROCEDURE

6.1. Overarching Requirements

1. All TEFCA Exchange leveraging Facilitated FHIR, including transactions that use out-of-band arrangements, as described in Section 6.3 of this SOP, MUST abide by the Common Agreement, QHIN Technical Framework, and Standard Operating Procedures (SOPs), except as stated herein.

2. All FHIR Adopters MUST follow the requirements as specified in the QHIN Technical Framework (QTF) Constraints Specific to Facilitated FHIR Exchange when initiating or Responding to a FHIR Query.
3. All use of certificates for HL7 SSRAA or other authentication must be consistent with the use of TEFCA certificate requirements as set out in the Technical Trust Requirements document.
    a. All FHIR Adopters MUST indicate in their Directory Entry the supported registration and authentication/authorization standards used and the FHIR Endpoint to be used for TEFCA FHIR exchange.
    b. Each FHIR Responding Node SHALL publish UDAP Server Metadata at `{baseURL}/.well-known/udap` for every `{baseURL}` published as a FHIR Endpoint in the RCE Directory.
    c. All JWTs used for completing UDAP flows according to the SSRAA IG SHALL be signed with a certificate corresponding to the TEFCA Facilitated FHIR Technical Trust Requirements policy.
        i. The x5c element of the JWT header SHALL contain the leaf certificate used to sign the request followed by the intermediate certificate which issued the leaf certificate.

6.2. FHIR Security

As of Q1 2027, all registration, authentication, and authorization of a FHIR client to a Responding Node MUST follow the requirements in HL7 SSRAA FHIR IG 1.1.0 – STU 1 US Sections 2, 3, 4, and 5. All FHIR Registrations SHALL use SSRAA dynamic registration for all clients. Specific milestones prior to this date include:

- Availability of testing systems with SSRAA support equivalent to all production Facilitated FHIR nodes
- All QHINs will have the capability to onboard and support SSRAA based authentication by November 1, 2026
    - Includes the ability to issue FHIR certificates and publish to the directory.
- All new FHIR Nodes published to the production directory after January 1, 2027 SHALL support SSRAA.

6.3. Exchange Partners

FHIR Adopters MAY determine their own exchange partners for the purpose of testing prior to 7/1/26 in Staging environments and 1/1/27 in Production. FHIR Adopters will be responsible for coordinating directly with each other to identify which FHIR Adopters have (1) a compatible registration and authorization approach.

This may result in FHIR Adopters not using FHIR Exchange with all other FHIR Adopters, which will not be considered a violation of the Framework Agreements, any applicable SOP, or the QTF.

6.4. Exchange Purposes

1. All transactions using HL7 FAST SSRAA Sections 4 and 5 Authentication and Authorization MUST use codes from the Exchange Purpose code system (OID: 2.16.840.1.113883.3.7204.1.5.2.1), as defined in the Exchange Purposes (XPs) SOP or an applicable Exchange Purpose (XP) Implementation SOP within the appropriate JWT and/or extension(s).
2. Notwithstanding the foregoing, if a QHIN has a TEFCA FHIR testing environment, they MUST test FHIR capabilities with any QHIN or Candidate QHIN who request non-production partner testing of such capabilities.

6.5. General Requirements

1. FHIR Adopters MUST share all patient data via US Core V6.1.0 conformant APIs, where such data exists and is available, consistent with applicable law. Adopters MAY make available resources conforming to versions higher than 6.1.0 or resources not included in US Core.
    a. All Responding Nodes with FHIR Capabilities listed in the RCE Directory Service MUST provide access to the Patient Resource and at least one additional US Core profiled patient compartment[4] FHIR Resource.
    b. In addition, all Responding Nodes MUST provide a CapabilityStatement resource and MUST use the FHIR CapabilityStatement resource to define FHIR server capabilities.
2. Any FHIR Implementation Guides MAY be used.

6.6. FHIR Endpoints & Endpoint Discovery

A required endpoint listing will be executed through a query to the QHIN's Directory following a Patient Discovery Query. These Queries will return FHIR Endpoints for locations where patient data exists. FHIR Endpoints returned in these Queries will not be limited to patient context and a patient search will be necessary to identify the specific patient.

1. All discovery of endpoints by Participants and Subparticipants MUST be executed by a Query to the QHIN Directory using the HCID returned from a Patient Discovery Query which will have the FHIR Endpoint(s) for Responding Nodes.
2. All FHIR-capable Responding Nodes MUST provide at least one publicly discoverable CapabilityStatement where CapabilityStatement.kind="instance".
3. All Responding Nodes with FHIR Capabilities listed in the RCE Directory Service MUST provide a CapabilityStatement for each endpoint associated with a FHIR server, defining the capabilities available at that endpoint.

[4] See https://www.hl7.org/fhir/r4/compartmentdefinition-patient.html for a complete definition

4. Capabilities listed within the CapabilityStatement MUST include all FHIR Implementation Guide operations supported by the Responding Node.
5. Capabilities listed within the CapabilityStatement SHOULD include all FHIR Implementation Guides supported by the Responding Node.

6.7. Patient Matching

Patient matching for the purpose of gaining or confirming the needed patient identifier for further data retrieval will conform to the FHIR Core Patient operation $match but use a US Core Patient Resource to allow for the additional demographics, including race and ethnicity.

1. All Nodes with FHIR Capabilities listed in the RCE Directory Service MUST support the FHIR $match operation using a US Core Patient Resource
2. $match operations MUST be executed using the demographics matching those used in the Patient Discovery Query as payload to allow for full Responses to Patient Queries from Initiating Nodes.
3. Responding Nodes SHOULD have the capability to return more than one potential patient match when a patient search yields more than one match.
4. Responding Nodes MUST NOT return more than one potential match when such action could be a violation of HIPAA or other Applicable Law.
5. When the Initiating Node specifies “onlyCertainMatches”=true within a $match Query Responding Nodes MUST honor that request by returning one and only one match, if a unique match can be found.
6. Responding Nodes MUST NOT return more than 100 potential matches when onlyCertainMatches is set to false.
7. All Initiating Nodes MUST include all available US Core Patient Resource demographics, including current and historical addresses, which can be sent and are not constrained by applicable law, within a $match Query for patient discovery with the exception of a Social Security Number, which MAY be included.
8. The Responding Node MAY fall back to requesting user specified credentials whenever authenticating with patient demographics per the Exchange Purpose (XP) Implementation SOP: Individual Access Services (IAS)fails or is not feasible.
9. Initiating Nodes MUST populate all Query elements in accordance with the adopted FHIR US Core's vocabulary bindings.

10. Initiating Nodes MUST normalize addresses to the Project US@[5] Technical Specification. However, if the field does not contain a street address but contains other geographical details, it is recommended that whatever information the patient provided not be abbreviated.
11. Demographics used in all $match Queries and Query Responses MUST follow all elements in the US Core Patient Resource, where available.
12. Responding Nodes MUST NOT require more than all US Core Patient Resource demographics before returning a patient list Response.

6.8. Provenance Use

The Provenance Resource will be used to track creation and transformation of data to and from FHIR resources. This will allow for accurate understanding of when a patient record is converted so that appropriate follow-ups can be made, where necessary. A FHIR US Core Provenance Resource MUST be available for Query following the FHIR US Core V6.1.0 Section 3.5 Basic Provenance Requirements and Section 13.149.1 Resource Profile: US Core Provenance Profile.

6.9. Error Responses

All error messages returned in Response to a FHIR Query will need to have sufficient information to allow for troubleshooting. This will follow the FHIR OperationOutcome Resource requirements and specification in FHIR Core R4.

1. Errors resulting from FHIR transactions SHOULD use the OperationOutcome Resource to return both human readable and machine processable information with sufficient detail to allow the client to determine if the error can be corrected at the client side.
2. QHINs, Participants, and Subparticipants MAY choose to obscure some of OperationOutcome details for security reasons. Any such choices SHOULD be linked to articulable security concerns.

6.10. Security

The following requirements are additional constraints on OAuth 2.0 [6] to further enable interoperability without reducing the security of transactions.

1. Registration, Authorization and Authentication for FHIR exchange MUST follow the requirements specified in section 0.
2. Authorization Servers SHOULD issue access tokens with a lifetime no longer than 60 minutes.

[5] *Project US@ Technical Specification.* – available at https://oncprojectracking.healthit.gov/wiki/pages/viewpage.action?pageId=180486153
[6] https://oauth.net/2/

3. An Authorization Server MAY issue a refresh token to an application using the Authorization Code Grant type if the Authorization Server issues a refresh token to an application that has requested and has been authorized to use the “offline_access”.
4. All implementations MUST support RS256, and SHOULD support ES256, ES384 and RS384.

6.11. Requirements for Use of HL7 SSRAA

Discovery

1. The SSRAA metadata endpoint MUST include the community parameter with the URI: urn:oid:2.16.840.1.113883.3.7204.1.5
2. The udap_certifications_supported metadata returned MUST include https://rce.sequoiaproject.org/udap/profiles/basic-app-certification.
3. The udap_certifications_required metadata returned MUST include https://rce.sequoiaproject.org/udap/profiles/basic-app-certification for all TEFCA registrations.

Registration

4. FHIR Initiating Nodes that initiate for multiple exchange purposes MUST register multiple client applications, one for each exchange purpose, to allow FHIR Responding Nodes’ Authorization Servers to provision them with the correct authorization policies for the given exchange purpose.
5. All client applications MUST register with an iss value corresponding to one of the client Subject Alternative Name (SAN) URIs present in its Facilitated FHIR Certificate.
    a. The SAN URI path MUST contain the exchange purpose for which the client application will conduct exchange.
6. The basic-app-certification certification JWT MUST contain the following elements:

certification_name	String with fixed value: "TEFCA Basic App Certification"
certification_uris	Fixed array with single string element: ["https://rce.sequoiaproject.org/udap/profiles/basic-app-certification "]
exchange_purposes	An array of one Exchange Purpose from the TEFCA Exchange Purposes SOP.
home_community_id	The HomeCommunityId of the Node making the registration request

7. If the Authorization Server returns a different client_id in the registration modification Response, the client application MUST use only the new client_id in all subsequent transactions with the Authorization Server.
8. If a new client_id has been issued for a registration modification, the responding Authorization Server MUST disable the old client_id so that it cannot be used for subsequent Queries.
9. Any retired client_ids MUST be preserved by the Authorization Server so that it can be associated with log entries and the Initiating Node.
10. If the client attempts to register for either a Client Credentials Grant or an Authorization Code Grant with a User scope but does not specify a user during registration, the server must respond with an "invalid scope" and not attempt to correct the scope to a System scope.

Business to Business (B2B) Client Credentials Grant

1. The authentication JWT extensions element MUST be present and contain a JSON Object containing the key "hl7-b2b" with a value equal to a B2B Authorization Extension Object.
2. Use of the hl7-b2b extension MUST conform to the requirements in **Error! Not a valid bookmark self-reference.**.

TABLE 2. TEFCA SPECIFIC HL7-B2B EXTENSION REQUIREMENTS

organization_id	required	RCE Directory Service entry Organization Resource id at the most granular level (e.g., lowest level Subparticipant or Child) formatted as a relative URI to the RCE Directory base URL (e.g., 'Organization/1.2.3')."
organization_name	required	String containing the Initiating Node's human readable organization name
subject_id	conditional	String containing a unique identifier for the requestor responsible for originating the Query. MUST be present when applicable
purpose_of_use	required	A length 1 array of strings containing the purpose for which the data is Queried, from the code set of authorized Exchange Purposes found in the Exchange Purposes SOP

3. When the B2B Authorization Extension object is included in a token request and the data holder determines that the authorization metadata submitted is insufficient for the data holder to grant access because the requestor has omitted the ACP parameter or has asserted a policy that is not acceptable to the data holder, then the Authorization Server MUST return an invalid_grant error Response to the token request, and this error Response SHOULD include the TEFCA-specific error extension as specified in **Error! Reference source not found.** in the 'extensions' object of the error Response.

TABLE 1. TEFCA AUTHORIZATION EXTENSION ERROR OBJECT

Extension Name: "hl7-b2b"		
Element	**Optionality**	**Requirement**
consent_required	required	The list of acceptable Access Consent Policy Identifier(s) corresponding to the asserted Access Policy required for authorization, an array of string values from the list of valid policy OIDs each expressed as a URI.
consent_form	optional	A URL as a string where the required consent form may be downloaded, if applicable.

4. Responders supporting use cases that require transmission of consent information MUST support the consent_policy and consent_reference claims and MUST be able to resolve a DocumentReference[7] or Consent[8] Resource included in the consent_reference array.

Individual Access Services (IAS) Queries

1. Responders MUST support the Authorization Code Grant type for IAS Queries.
2. The responder MUST support the TEFCA IAS Authorization Extension Object identified by the extension key "tefca_ias" as defined in **Error! Not a valid bookmark self-reference.**.
3. A client application requesting a token for patient requests using Client Credentials Grant MUST include the TEFCA IAS Authorization Extension Object in its token request in addition to the hl7-b2b Authorization Extension object during the authorization flow.

TABLE 4. TEFCA IAS AUTHORIZATION EXTENSION OBJECT

Extension Name: "tefca_ias"		
Element	**Optionality**	**Requirement**
version	Required	Fixed string value: "1"
user_information	Required	FHIR RelatedPerson Resource with all known demographics. Where the user is the patient, the value of the relationship element MUST be "ONESELF"
patient_information	Required	FHIR US Core Patient Resource with all known demographics
id_token	Required	The CSP-provided OpenID Connect token as further defined in the Exchange Purpose (XP) Implementation SOP: Individual Access Services (IAS). Responding server SHOULD respond with invalid_grant if missing.

[7] See https://hl7.org/fhir/R4/documentreference.html
[8] See https://hl7.org/fhir/R4/consent.html

Extension Name: "tefca_ias"		
Element	**Optionality**	**Requirement**
consent_policy	Conditional	If user is not the patient, this element represents the Access Consent Policy Identifier corresponding to the asserted Access Policy that represents the identity proofing level of assurance of the user, array of string values from the subset of valid policy OIDs in**Error! Reference source not found.** that represent identity proofing levels of assurance, each expressed as a URI, e.g. ["urn:oid: 2.16.840.1.113883.3.7204.1.1.1.1.2.1"]
consent_reference	Optional	An array of FHIR Document Reference or Consent Resources where the supporting access consent documentation can be retrieved, each expressed as an absolute URL, e.g. ["https://tefca.example.com/fhir/R4/DocumentReference/consent-6461766570"]

4. The user metadata submitted by the requesting application in the patient_information element of the TEFCA IAS Authorization Extension Object MUST correspond to the verified identity attributes of the permitted user who is making the Query.
5. If the submitted user information does not sufficiently match a person known to the responder, or if the responder does not support this workflow for IAS Queries, it MUST return an invalid_grant error in Response to the token request.

6.12. Scope Negotiation Requirements

FHIR Server scope negotiation MUST conform to following constraints:

1. The `scopes_supported` metadata **SHALL** be present in the .well-known/udap object and **SHALL** list all supported wildcard scopes.
2. Client applications and servers **MAY** support wildcard scopes.
3. A client application **MAY** request a wildcard scope only if wildcards are specified in the server's scopes_supported metadata list.
4. If a client application requests a wildcard scope and the server supports wildcards, then the server **SHOULD** return either the wildcard scope or an expanded set of scopes that the client has been granted in its response.
5. If a client application requests a wildcard scope and the server does not support wildcard scopes, then the server **SHOULD** respond with an error of "invalid_scope" or "invalid_client_metadata", as appropriate.

6. If a server supports OIDC or SMART App Launch scopes, the server **SHOULD** put the corresponding scopes (e.g. “openid”, “offline_access”, “email”, “fhirUser”, etc.) in their `scopes_supported` metadata.
7. A server **MAY** grant fewer scopes than requested by the client application if the client application cannot have a scope specified in the request based on technical or policy guidelines at the responding organization or if the server does not recognize one or more of the requested scopes.
8. A server **SHOULD** respond with “invalid_scope” or "invalid_client_metadata", as appropriate, only if a wildcard scope is requested and not supported, or if none of the requested scopes are supported.
9. At the time of a token request, an authorization server **MAY** grant additional scopes that are not in the set of scopes requested by the client application if the application has been registered with the server with a different set of scopes than was requested at registration based on technical or policy guidelines at the responding organization.
10. Scopes granted by a server to a client application at the time of an access token request may be the same as the set from registration or be a subset.
11. Scopes granted by a server to a client application at the time of an access token request may be the same as the set of scopes requested by the client application or be a subset.
12. An application **SHOULD** be able to receive a superset of the scopes requested if the server’s policies dictate that a request with a certain system or user/user role is granted specific scopes that are not part of the original request.
13. A server **SHOULD** return “invalid_scope” or "invalid_client_metadata", as appropriate, only if none of the scopes requested are available and/or not part of the scopes requested during registration.
14. A server **SHALL** include the `scope` parameter in a token response if the set of scopes granted by the server to the client application is not identical to the set of scopes requested by the client application, or if the client application does not include a set of requested scopes in its request.

7 VERSION HISTORY

Version	Revision Date	Section #(s) of Update
Version 1.0	Released July 1, 2024	N/A
Version 2.0	February 6, 2026	Sections 4, 5, and 6