

Mariann Yeager
Sequoia Project
8300 Boone Blvd Suite 500
Vienna, VA 22182

September 17, 2021

Re: QHIN Technical Framework Draft v1

Dear Ms. Yeager,

On behalf of CommonWell Health Alliance, we are pleased to submit comments on the RCEs QHIN Technical Framework Draft v1 as published July 28, 2021. We have always supported a fully connected nationwide ecosystem that links providers of all types and recognize we cannot do this alone. We fully support the goals and objectives of TEFCA and look forward to being one of, if not the first QHIN recognized in the country.

CommonWell Health Alliance is a not-for-profit trade association made of various health IT and health care stakeholders. As a membership-based trade association, we provide an environment to openly work on interoperability improvements across many cornerstones of health care including but not limited to technology companies, payers, State and Federal agencies, providers, clearing houses and patients. When the Alliance launched eight years ago, we started with services centered around Care Treatment and provided the ability for providers to query across other provider systems and retrieve data about a given patient. We have been and continue to be a patient-centered network available nationwide and are proud to have added Patient Access use cases to give individuals the ability to find and access their data through patient portals, personal health records and other patient-centric applications.

CommonWell has a simple vision: health data should be available to individuals and caregivers regardless of where care occurs. Additionally, access to this data must be built into health IT at a reasonable cost for use by a broad range of health care providers and the people they serve. At CommonWell, together with our service provider and members, we have created and deployed a vendor-neutral platform that breaks down the technological and process barriers that inhibit effective health data exchange. We leverage existing standards and policies in order to enable scalable, secure and reliable interoperability as easily as possible for our members and their customers across the nation. We are excited to see TEFCA is formed with a similar use of standards and think it has the potential to create the nationwide system we dreamed of through its interconnected system of Qualified HINs.

While we fully endorse the fundamentals of TEFCA and what the technical framework outlines in QTF Draft v1, we do have five specific areas for comment. The first three center around a potential challenge regarding the ability to build and maintain trust across the connected TEFCA system. Trusted is the first word in the acronym TEFCA; any potential for erosion of trust in the initial setup of TEFCA is a concern. The last two areas of comment center around focus – simply put, making sure we work towards creating a workable solution within the current timeline as outlined by the ONC without distraction. These five areas are detailed further in the remainder of this document:

- Area of Concern #1: Elimination of requirement for QHIN to have an RLS or eMPI
- Area of Concern #2: Patient Matching Standard(s)
- Area of Concern #3: Transforming Documents to C-CDA v2.1
- Area of Concern #4: Message Delivery – XCDR for QHIN-to-QHIN push
- Area of Concern #5: FHIR Roadmap

Area of Concern #1: Elimination of requirement for QHIN to have an RLS or eMPI

Reference: Page 11 of the Draft QTF

Each QHIN has either a Record Locator Service (RLS) OR Enterprise Master Patient Index (eMPI) OR the ability to query all of its Participants for a patient lookup within the timeout limitation as specified in the QHIN Service Level Requirements Policy

As the first transaction in the document query chain, XCPD is the most critical of the three. If a system or endpoint is unable to respond in a timely, accurate manner to patient match transactions, the subsequent XCA-Query and XCA-Retrieve will never occur. This means it is critical that XCPD responders are highly optimized and prepared for a level of scale beyond what any of us have experienced so far.

The foundational goal of TEFCA was the ability to query everyone connected across the TEFCA ecosystem for any patient at any time via any TEFCA on ramp. CommonWell was founded on this identical principal (ability to query everywhere, every time), and we can accomplish this inside our network because we know where records are located for every patient query before the query arrives; in short, we have an eMPI-driven Record Locator.

If the “OR” statement underlined above remains, it might technically work, but it is likely to be slower and less reliable than a focused eMPI/RLS implementation. Assuming this option remains, we believe it will be absolutely critical to develop and enforce a carefully defined SLA for XCPD performance. We believe this SLA will need to include flow down guidance to participants behind the QHIN so they understand their performance obligations to make TEFCA fully operational. Without flow down, there will likely be concern that some QHINs may be missing potential responses from their participants.

None of us wants to have endpoints consistently or intermittently not responding, particularly if it is difficult to understand who didn't respond when. A lack of transparency would serve to erode trust in TEFCA.

Scalability

The original TEFCA design philosophy did not assume endpoints behind the QHIN needed to be able to handle the same level of scale as the QHIN itself. With this RLS/MPI-less option, this philosophy is fundamentally broken. The millions of XCPD transactions likely to hit each QHIN would also hit the endpoints behind it. We believe there are circumstances where this fan-out methodology could work, but they are likely very limited. We also have concerns that less

mature QHIN candidates attempting to use XCPD fan-out for patient discovery will be overly optimistic in their ability to build an RLS/MPI-less option, and we will not know if it is not working until very late in the game.

As the QTF indicates, the solution to this problem could be a performance SLA. We agree this is needed at a minimum, but given our concerns above, we encourage this to be a tight, fast response SLA with flow-down terms to participants. These flow-downs could be contained in the SLA itself, or the RCE could require they exist in QHIN agreements with their participants with review for sufficiency as part of the QHIN certification process.

For a point of reference, the CommonWell SLA for patient discovery is currently under 2 seconds for transactions inside our network. Outside our network (e.g., Carequality), our performance goal is not as tightly defined nor is there an SLA required by the Carequality framework.

Regardless as to our current external performance specificity, we support a tight SLA for all XCPDs TEFCAs-wide and welcome it being applied to us along with everyone else.

Example of Erosion of Trust

TEFCA is designed to create trust across the system. This not only includes trust in privacy and security but also in technical capabilities. Take this scenario:

- The XCPD QHIN-QHIN response SLA is 3 seconds maximum
- There is a QHIN with 10,000 entities behind it, and it does not have an RLS/MPI
- An incoming ITI-55 hits a QHIN who has implemented - Alternate Flow 3 (XCPD fan-out), as described on pages 12-13 of the QTF Draft
- 3 of the QHIN's entities have data for the requested patient
- 1,000 of the entities respond in a timely manner (1 second) to the responding QHIN
 - These 1,000 do not have the 3 patient records
- 9,000 are slower to respond to the QHIN, with an average response of 4 seconds and none faster than 2 seconds
- The QHIN closes all internal XCPD connections in 2 seconds so it can package responses to make the 3 second SLA
- The result is the QHIN responds with "no patient found" to the requesting QHIN

While the above may sound like extreme numbers, even at lower numbers of non-responders due to timeout, trust in the ecosystem will be negatively impacted.

The technical root of the problem stems from ITI-55 not having an error protocol to represent failed transactions from multiple entities behind the XCPD responder. An XCPD response can indicate the response is a partial response, but the user would not get an indication of how partial it is. If there are 10,000 entities behind the QHIN, and the 9,000 with this patient's records could not respond in time to meet the SLA, under normal ITI-55 operations there would be one "patient not found" message with an optional indicator available to indicate that the results may be incomplete. Given this, how would a querying party know they aren't getting all the data and have visibility to where they may be missing content from?

To our knowledge, the current standard does not support this well, so we would need to invent a new standard or an extension to the existing standard, which takes time. Even if/when we figure this it out, it will likely only serve to highlight the issue, not fix it.

Understanding the Scale

Carequality is the broadest framework today for exchange. CommonWell is one of the biggest Implementers connected to the Carequality framework. Not all CommonWell Service Adopters use our broker to exchange with Carequality entities; our second and third largest Service Adopters belong to and have implemented both Carequality and CommonWell. Even so, the scale of exchange we see daily is massive and continues to grow.

In the last measurement month, we had more than 180 million outbound XCPD transactions and responded to more than 90 million XCPD queries per week; that is Carequality alone and does not include the volume contained within our own network. Also, it is worth mentioning that Carequality is focused on directed query not broadcast. Most entities in Carequality have a geozip search algorithm to make an educated guess as to where the data holders are for a given patient. Some have other data sources to help inform where to search, and some have more advanced algorithms than others; but in the end, many Carequality discovery transactions yield no matching records.

Today, when we query out to Carequality participants, our observed patient discovery match rate is about 10%. When Carequality participants query into CommonWell, the match rate is about 4%. We actually don't know if the different rate of matching is because our directed query algorithm is better at targeting possible data holders or our patient matching response algorithm is more restrictive (though this is unlikely because we also filter responses back from Carequality endpoints through the same matching logic we use when responding). What we do know is that it is a lot of volume, that it is growing, and that it is likely to go up even more when TEFCAs are in full swing.

Within Carequality, all Implementers perform directed queries; entities are not allowed to broadcast a patient match request to every endpoint in the framework. Within TEFCA, a broadcast query to all endpoints facilitated by the QHIN is the expected, normative flow. If a QHIN does not have an RLS and fans-out incoming XCPD to its underlying entities, this could mean thousands of endpoints could receive millions of transactions a day. It is architecturally simpler to scale one system for millions of transactions than to scale hundreds or thousands of entities for the same load.

This is precisely why CommonWell was built with an RLS capability from day one. We recognize there are vendors and customers with the resources to handle the high volume necessary to make a fan-out XCPD and an associated QHIN-QHIN XCPD SLA work; but that is not the norm. We are concerned that some may not understand the potential scale and will only find they cannot meet the SLAs after we collectively reach production volume levels. This could create an undesirable allocation of haves and have nots, and/or an erosion of trust through lower performance and missing data TEFCA-wide.

Patient Matching Heterogeneity

Potentially amplifying the trust problem is the fact that the QTF does not currently contain a standard for patient matching. While we think a patient matching standard should be incorporated (more on that below), we would at least expect there to be a process or expectation to get us there. An XCPD fan-out patient discovery methodology makes this much more difficult to develop later because there would be more parties that need to come to a consensus and implement it.

Let's say there are 15 QHINs. If all had an RLS-based system, we would only need to develop and implement a standard algorithm with 15 entities. But, if one of those QHINs uses the fan-out XCPD methodology and it has 10,000 participants, we now have 10,014 entities that need to deploy the agreed to logic.

We should not be blind to the realities of software development and deployment. Rolling a change out across more vendors and more customers makes it harder for us to raise the bar. To be fair, this problem is likely less if the QHIN is a single-vendor QHIN with a cloud deployed EHR, where one software change updates all 10,000 of their customers; however, assuming patient matching is a core technical function of TEFCA, we should be careful to avoid this known pitfall upfront.

Relation to Patient Discovery Alternate Flow #1

Finally, we have some concern with Alternate Flow #1 in relation to trust erosion from uneven performance under Alternate Flow #3. The QTF draft described XCPD Alternate Flow #1 on page 12 as:

The following flow may be used when the Query Source only wants to query one or more specific organizations

This is very similar to the method of exchange under Carequality. The difference here is Carequality has operated on the general guidance that broadcast queries should be avoided, while this alternate flow does not define an upper boundary of usage and could result in this happening.

In the event a QHIN or participant believes there are entities behind another QHIN not responding in time under Alternate Flow #3, they could resort to mass usage of Alternate Flow #1 to better understand what is occurring. Using the flow would allow the querying party to directly observe the responding party's match/no-match/timeout response from each individual participant to help them determine what may be missing.

We believe that directed queries should be allowed and absolutely have a place. We also believe falling back to mass amounts of transactions following Alternate Flow #1 is not extremely likely and is not an expected result of the design of Alternate Flow #3. Regardless, we believe Alternate Flow #1 should have a defined boundary condition(s).

Recommendation if the RLS/MPI-less Option Remains

- Tight SLAs with performance guarantees and associated penalties
- Include review of SLAs as part of the QHIN certification process
- Expand implementation document to educate potential QHINs on operational expectations particularly with regard to patient matching
- Define boundary conditions for usage of XCPD Alternate Flow #1 to avoid potential broadcast misuse

Area of Concern #2: Patient Matching Standard(s)

Patient matching across systems is key to TEFCA; however, the QTF has no mention of a standard approach. Public QTF information sessions have stated there is no intention to include a patient matching standard because this is out of scope for TEFCA. Quoting one CommonWell Member in a recent discussion on this matter: “If patient matching standards are out of scope for TEFCA, when will it ever be in scope? This feels like the most appropriate time and place ever to lock down such a standard.”

First, let us recognize that the QTF does reference USPS Publication 28 as a standard for documenting patient address(es). This is a good first step, but for this to be truly effective, we need to explain how the standardized address fields should be used in the matching process across systems.

Through our participation in the Carequality framework, which does not have a matching standard aside from a minimum set of fields you must include in an XCPD, we have seen evidence that the lack of a standard method for patient matching has been an impediment to trusted exchange. Some Implementers use deterministic algorithms, others use probabilistic algorithms, and for many others we simply don't know what they are using.

Matching Standard for Patient Right-to-Access Requests

This problem will be felt the most if TEFCA requires all QHINs and their participants to respond to patient right-to-access requests across QHINs. We will likely discover that TEFCA does not work for patient access in practice without a patient matching standard. Homogenizing the match logic across all Use Cases may not be necessary, but we feel it is necessary to have a standard for patient right-to-access or this use case is unlikely to be adopted.

The ONC does have the authority to set a standard algorithm or process, and we encourage the RCE to work with the ONC to do so in the context of TEFCA today.

We believe if there was a specified patient match algorithm and the ONC stated that any algorithm more restrictive than the defined standard is “unreasonable”, then TEFCA and the goal of ubiquitous patient access would be aligned. This defined algorithm would become a safe harbor of sorts for participants who could then point to this standard in a statement like this:

“We are matching in a manner consistent with the technical requirements and standards set forth by the ONC, the federal agency responsible for developing and enforcing standards for health IT in the United States”.

Deterministic Matching

One suggested first step would be to define deterministic matching and require that no one can match with a more restrictive logic across any use case. Across the industry, we regularly discuss patient matching, including deterministic methods. The reality is we don't have agreement on what fields, when used together, make a match deterministic:

Does deterministic mean a perfect string match for every demographic field? Can you assume street address synonyms like Unit, Apartment or Condo are functionally equivalents? Or is it really the full address line as typed? Are you allowed to clean up an address and normalize to USPS-designated addresses before you try to match a person?

Many companies and providers would likely agree that given name, family name, gender, date of birth and zip code are not quite enough to be deterministic. But what if we include the first line of someone's address? Previous papers, including one published from the Sequoia Project (available to download here: <https://sequoiaproject.org/resources/patient-matching/>), have shown that address is a highly available, relatively distinct field; when used with other attributes, it can be used well to create good, low false positive matches. But, as that Sequoia paper also points out, the address line is traditionally hard to use because it is a non-coded string.

In CommonWell, we found that if this string is normalized to the official address recognized by USPS as the correct address (the address as assigned by the local municipality for the address), it can be used quite effectively for cross-entity patient matching. We normalize addresses (e.g., clean out extraneous characters, normalize apartment/unit/suite/slip naming to the modifier issued by the local municipality, fix zip codes to appropriate 5+4 codes for an address, etc.), but we will only auto-match two records if it follows our version of deterministic.

The fields we use for auto-matching are:

- First name
- Last name
- Date of birth
- Gender
- City
- State
- Street Address
 - Alternative to street address is a phone number or e-mail address match

When necessary, we have multiple matching logics and sometimes use probabilistic algorithms. For example, matching outside our network (e.g., requests from Carequality sites) is probabilistic. In Carequality, we only exchange for TREATMENT, and the match threshold for provider-to-provider can generally be lower than for provider-to-patient due to the mutual trust embedded in the framework for TREATMENT.

Provider-to-Patient Matching

We believe that having a defined algorithm, using it universally, and defining anything beyond this boundary as unreasonable could eliminate the optionality of responding to patient request via query-based document exchange. If TEFCA requires all participating endpoints to respond to a patient REQUEST, a defined match standard could make the voluntary adoption of TEFCA more palpable.

Today, many requests for patient access are blocked because of security concerns related to potential incorrect matching. If the entire market is working with the same reasonable standard as defined by the ONC, we believe the use of patient matching as a blocker (an arguably necessary one for now) would dissolve.

As a side benefit, this could empower patients to fix missing data issues on their own.

Knowing the implementation of a national health identifier is not likely to happen anytime soon, we need to have a conversation about best practices and implementation expectations with what we have to work with, and it needs to happen now. TEFCA is likely the best place to have and use the results of that conversation in practice.

Recommendations

- Leverage the ONC's authority to set a patient-matching standard
- Maintain patient access as a required purpose of use under TEFCA
- Develop educational content describing the record location and matching process, how to determine which providers are participating in the TEFCA network, and how the patient can correct any suspected errors in their demographic data in addition to their clinical content

Area of Concern #3: Transforming Documents to C-CDA v2.1

Reference:

QTF-039: If a Document Retrieve response is not in C-CDA 2.1 format, QHINs MUST convert the response to C-CDA 2.1 format except where consistent with QTF-043 and QTF-040.

First, this is currently infeasible for many candidate QHINs. While some QHIN candidates receive data from their participants and sign XCA documents' responses as their own, many are fully federated with documents signed by the individual organization. As such, many QHINs will not have the right to open documents as they pass through their systems, making it infeasible for the QHIN to convert from one document type to another.

Second, this is too prescriptive as an absolute requirement in the QTF. There are many documents passed today that are of high value to recipients but don't follow the C-CDA 2.1 format (e.g., older CCR, CCD or CDA C-CDA R1). Likewise, it would be expected that future document and content types may not be in this format, particularly as new electronic formats are developed.

Third, conversion takes time and resources. If this requirement remains in the QTF, new SLAs must be created, as this is not a process currently performed by many QHIN candidates.

Fourth, converting documents may cause unexpected, potentially clinically dangerous changes to a responder's data. Other document types may not contain the necessary content to be converted due to different vocabulary standards, changes in documentation norms over time, etc. This could yield to mistranslations or omissions of data, all without the knowledge of the data requestor or the data provider.

Recommendations

- Eliminate QTF-039
- Leave conversion, deduplication and other transform services as examples of potential value-add processes QHINs may offer to their own participants
- Focus on ensuring new documents are generated by the source in accordance with latest standards (e.g., those specified in ONC's certification program)

Area of Concern #4: Message Delivery – XCDR for QHIN-to-QHIN push

The QTF requested commentors focus on three possible options for push in the QTF: 1) Require, 2) Defer or 3) Include.

CommonWell Health Alliance elects option 2: Defer “QHIN Message Delivery” from QTF until a FHIR-based solution is readily available.

We support push and think it is an important partner to query-based exchange, but the selection of XCDR as the technical standard for push does not fit into TEFCA's build philosophy of working with mature standards with implementation specificity. The reality is XCDR standard is not widely adopted by the industry nor by candidate QHINs at their edges or on the inside of their networks. The adoption of XCDR would, therefore, take significant effort by QHINs, QHIN participants and sub-participants when there are other methods of push that are more widely deployed.

In discussion on QHIN push, we have heard comments indicating XCDR would be limited to QHIN-to-QHIN exchange, not what happens inside the QHIN. Even if you limit XCDR to the edge of the QHIN (for QHIN-to-QHIN exchange), any push method inside the QHIN is still likely to be a significant effort. Participants behind the QHIN would need to build XCDR, or the QHIN will need to develop bridge technology to translate XCDR to whatever is used by their participants. The reality is, many candidate QHIN entities do not have robust push capability beyond moving ADTs or basic event messages and are not structured as intermediaries for push-mode exchange.

FHIR Message Delivery

Regarding FHIR as a target for future push development, FHIR-push protocols are not sufficiently mature nor adopted either, which again breaks the mature standard build philosophy of TEFCA. Direct, for example, has broader adoption and has been effective in enabling general push workflow as well as specific ones (e.g., the case reporting push flow for APHL under a trusted exchange framework approach across CommonWell and eHealth Exchange with APHL under the Carequality agreements). FHIR push, using a combination of writes, messaging and subscription methods, is showing promise and is likely to be implemented in EHRs in the next couple of years, so it does deserve some attention. That said, it is not widely believed that the initial slate of FHIR-push use cases would require intermediaries at all. There are some workgroups developing use cases where intermediation may be helpful, but these are very new.

Intermediary Push Transmission

If QHINs are not intermediaries for push transmission, i.e., not required to broker message delivery, the next logical place for TEFCA entities to participate is in the establishment of a robust directory that should include the push capabilities of all participants for whatever method(s) agree to be included in the directory within the TEFCA framework: Direct, XDR, XCDR, FHIR etc. There are likely many more use cases that would benefit in the short (and likely long) term from a trust framework with reliable endpoint targeting instructions supported by a scalable directory infrastructure. In this model, the QHIN is not directly involved in the transaction, rather it extends the framework's policies and obligations to enable technologies to operate at scale through the reduction of point-to-point data sharing agreements while improving patient and record location with endpoint/address discovery.

Recommendations

- Defer Push Message Delivery without a specific implementation target
- Hold an open discussion on the need for intermediation and what parts of the technology path may need it
- Collaborate with FAST and related workgroups to develop an industry roadmap for push transactions – both brokered and non-brokered
- The future governing committees of TEFCA should work with the ONC and the RCE to decide if push belongs in the TEFCA roadmap, and if so, in what capacity

Area of Concern #5: FHIR Roadmap

We believe in FHIR as a future method of exchange; however, it should not be explicitly defined in the context of a TEFCA roadmap yet. The timeline to operationalize TEFCA is tight, and while it leverages many mature standards and capabilities, many candidate QHINs will need to build and operationalize new capabilities. We believe TEFCA v1 should focus on query-based document exchange. Developing the FHIR roadmap too early is likely to be an embedded distraction. It may also cause some participants and candidate QHINs to defer their participation, thereby not getting us to the scale we want out of the gate. We support working on this in parallel but think it is too early to be prescriptive in the QTF and/or the Common Agreement.

Recommendations

- Continue to acknowledge the importance of FHIR and that its notable absence in the QTF is intentional
- Educate the market on the TEFCA update process, which includes the ability to add, change, supplement and substitute modalities of exchange as market needs and capabilities change

Note on Our Comments

These comments are reflective of the opinions of the Alliance and its members in regard to the objectives of CommonWell. It is not intended to represent the individual comments of each of our Members. Nor are these comments made to represent the view of any particular member. We fully expect some of our members and their customers to submit their own comments individually.

Final Comments

The Alliance remains committed to patient-centric interoperability on a national scale with the goal of ubiquitous, secure exchange of clinical data to the benefit of providers, patients, payers, public health and other stakeholders in health care and beyond. We are encouraged by the mission and vision of the ONC and affirm our commitment to being a strong partner in executing on the vision.

On behalf of CommonWell Health Alliance, thank you for the opportunity to comment on the latest QTF Draft. For any clarification or comments, please feel free to contact me at paul@commonwellalliance.org.

Sincerely,



Paul L Wilder

Executive Director
paul@commonwellalliance.org

CommonWell Health Alliance
75 Arlington Street, Suite 500
Boston, MA 02116