16870 West Bernardo Rd. #400
San Diego, CA 92127
June 17, 2019

Donald Rucker, M.D.
National Coordinator for Health Information Technology
U.S. Department of Health and Human Services
330 C Street SW, 7th Floor
Washington, D.C. 20201

Dear Dr. Rucker,

Thank you for the opportunity to comment on the TEFCA NPRM. Please find below our comments on the second draft.

Sincerely,
Julie Maas
EMR Direct

General comments

Provide a single "on-ramp" to nationwide connectivity:

This implies reusable identities. Whatever the transport methods used within a Common Agreement, why not take existing models that work as a starting point? In other words, leave existing push and pull frameworks where they are, and build from there, so as not to interfere or threaten what is already working.

Requiring that a reusable credential could be provisioned and allowed to connect to any HIE would potentially be a more tractable problem than requiring every HIE (plus all other QHINs) to develop and cross-test the capability of sending out and then answering any targeted or broadcast query to any/every other HIE (plus the other QHINs). The HIE network to network model on its own would present plenty of interoperability challenges and cross platform compatibility issues in and of itself.

The TEF doesn't address the bigger issue of payload incompatibilities. The CCDA scorecard is a great starting point, but there is still a significant issue relating to these summaries not being able to translate from one institution to the other. This is the interoperability "elephant in the room". ONC said a few years back that you wanted to go beyond "throwing the CCDA over the wall" so this will be an area to improve. If we don't, those billions and billions of CCDAs getting sent over Direct or query or FHIR and the improved interoperability you're trying to produce won't matter. If a CCDA arrives in EMR B from

EMR A, how often are they able to use and to incorporate it? Successful patient match is part, but ONLY part, of this problem. Cross-testing is challenging to coordinate.

With regard to push use cases, the Direct community includes approximately 2 million endpoints using low-cost, point to point, accredited network services and approximately 30 million messages are moved each month with this infrastructure. That's more than query, and growing all the time. The absence of support for this network from ONC in the draft TEF is a bit disappointing, to say the least, especially considering that it was ONC who required that all certified Health IT build the capability into their systems. It's obvious that Direct is a threat to less scalable (and more lucrative) business models. Although it isn't as highly revenue generating, it certainly gets the job done. We should be sure to use it when we can, and why not use it as it has been implemented, rather than increasing the complexity and cost of interoperability by creating yet another variation of it?

A Direct service address would lend itself very well to cases in which a sending party wants to push data to another system but does not have an originating Direct address of their own. Such service addresses could be provided by HISPs to a QHIN needing to transmit to a Direct endpoint. Since every provider utilizing certified Health IT is supposed to have a Direct endpoint, it seems like the service address construct would only be needed for payers or other entities not yet required to have one, or for senders requiring the assistance of the QHIN mostly for its directory capabilities. A similar construct can exist on the receiving side, where message elements can indicate the provider for whom a message is intended, if the recipient's organization name as discoverable via directory service is not sufficiently granular or if the entity receives via a service address managed by their QHIN.

Enable Electronic Health Information to securely follow the patient when and where it is needed:

This is a great ambition but we have some work to do to define who the patient is. When new ONC rulemaking relating to patient attributes and their verification is established—including the addition of verified Driver's License and passport numbers, healthcare plan numbers, the new medicare identifier number, physical addresses that have been validated using usps.gov or even better verified using credit bureau type records as belonging to a particular individual, and Direct addresses, email addresses, OpenID identifiers, or other electronic attributes that can be verified using Userinfo Endpoint, demonstration of control through email verification codes, or mail to the home address (which would be a great step-up from LoA-3 to IAL2…the California DMV does this as a 2nd proof of residency when provisioning a REAL ID)—vastly improved matching should result.

It is important not just to enable identity verification, but to be able to make sure the person requesting Individual Access Services is also the same data subject so verified. Today, the proofing requirements for patient portal access are not as rigorous as the TEF requires, but there are numerous controls in place that keep my data requiring sub-LoA-3 proofing to only be appropriately accessed. The issue seems more about institutions that provide no reasonable path toward electronic access, not a solution that patients should feel better about just because it is IAL2. Though it is a natural progression from LoA-3, IAL2 still has issues to be worked out—among them the lack of driver's license verification services in every state. Suggest starting with best practices for what is working well in the field today (verification using credit bureau type records), rather than a blanket endorsement of IAL2. The 800-63-3A document doesn't even make it clear whether notice is required in either in-person or remote cases. Mailing something to the verified address is a key step that increases confidence and gives people awareness that their identity is being used in LoA-3 processes today. If careful IAL2 proofing is intended to be

stronger than LoA-3, it will be sure to include notice mailed to the physical home address of record or demonstration of control of an electronic address that is linked to the individual in authoritative records appropriate to the use case. Additionally, remote proofing with IAL2 seems un-attainable for those without either access to data from mobile network operators (to confirm that a mobile number is registered to the person whose identity is being verified) or some pre-existing relationship with the person being proofed.

Support nationwide scalability:

Something that seems to be missing here is an overarching architectural design. Perhaps it is out of scope for these documents, but what is the ONC envisioning as the long term role for HIEs? If some information is available through FHIR servers and some is available straight out of the EMR or through my PHRs that may participate in Carequality, how does the TEF prevent any new ER-initiated query from pulling back too much data—due to redundancies in such a networked system? TEF task force members have also envisioned such a data duplication issue. It is possible at a smaller scale for FHIR clients to generate repositories of duplicated data also, if data identifiers are not used carefully by FHIR data holders and the clients that query them. It's concerning to attempt to envision about how this can be done safely without a lot of improvement in patient matching that includes care to avoid unintentional replication, as well as standardization and adoption of industry-wide best practices.

"Network of networks":

In the most practical sense, it seems this would be implemented using an OpenID identifier that is capable of containing embedded identity information and that can be used to authenticate either to a HISP service for message transmit or receive, or to an app that can use a hosted, trusted certificate to sign assertions which are also bound to a Participant Member's identity.

Individual Users:

This nomenclature seems potentially confusing and a better title for this entity would be something like Data Subject. It would be great if all Data Subjects also had a role in the Common Agreement and could track the use of their data across QHINs—even managing their Meaningful Choice more directly.

Relationship to HIPAA:

It is a known limitation of the app ecosystem that apps downloaded to a phone are incapable of preserving their identity on the network whereas web-hosted applications are. Without sacrificing the ability of patients to use such apps should they wish to do so, an industry effort to recognize and indicate the added confidence in the use of such "confidential clients" (and in fact the ability to restrict use of public clients in a professional setting—e.g. broadcast query, for one) is an operational detail that is likely out of scope at this level but worth noting well in advance so that this variety app types can be managed from the start.

Agreement Structure:

This seems so burdensome. Might it work for cross-organizational FHIR queries, and for a QHIN of HIE(s), for which there are not yet any broad based common agreements, to leave other networks' Participant Member Agreements as they are but layer on a "network of networks" addendum for those who wish to step up to Common Agreement transactions?

Additional Data Exchange Agreements Permitted:

This is fine to say, but a network of networks, if not carefully designed to complement entrenched, though still in high-growth, interop solutions that our citizens and providers have paid high tolls to build, will cannibalize those existing networks if their progress is not allowed to experience sustained, uninterrupted growth. It is important not to introduce a program that is akin to putting everyone on AOL but without the ability for their AOL email addresses to interoperate with any other email address. We would simply be subsidizing the use of AOL and diluting interoperability previously achieved, even if we don't explicitly prohibit the use of alternative solutions.

Meaningful Choice and Written Privacy Summary:

Saying that we want to do this, and enforcing it, requires a more meaningful representation of patient identity across the care continuum, or at least in as much as the patient's data is accessed through the Common Agreement. That is to say, in order to prevent sharing that is inconsistent with Meaningful Choice, data that IS shared will need to map to a unique patient in order to be certain we are operating reasonably (that data for patients who have opted out is not being exchanged). A national patient identifier, or a proxy for the same that uses credit bureau type records to verify underlying attributes and establish uniqueness, will need to be an integral part of this system. It would be optimal from a patient perspective for Meaningful Choice to be defined more broadly, such that it might have impact where data can be purchased and in other data sharing agreements outside the TEF.

Appropriate Security Controls:

This is a good topic for further investigation and merits a risk analysis to evaluate what controls may be necessary. The statement itself makes it clear that there may be some archival of records passed from a participant to a participant member as part of supporting a query, and governance around what data may be persisted and for what purpose is likely also needed. For example, as a query facilitator it may not be appropriate for a QHIN to archive all the data passed through it, unless it offers a caching service or archival service and this is clearly communicated to participant members and to endpoints themselves (such use may be considered an extension of the permitted purpose a Participant Member may be asserting) or in case metadata is collected for the purpose of logging and helping to build the necessary indexing of a patient's records which will become necessary in order to access all records for a given patient.

No Use or Disclosure Outside the United States:

It seems like the QHINs might be better informed of circumstances requiring such a query through the use of credentials explaining the need (for example, certificates issued to embassies or medical centers located outside the United States but who are authorized to participate in the Common Agreement will have their intended location reflected in their certificate so that it would not present a concern to the policy engine of a QHIN receiving the query that the request originated outside the US—the Participant Member was vetted in the context of expected international use).

MRTCs Overview:

It seems like it would also be worth ironing out these areas that seem to impede interoperability since existing networks will likely continue to be used in their interoperability-impeding form. Until such time

that the country can afford to use the different technologies with their new onboarding requirements in the impedance-preventing design being prescribed, will the TEF also require entities that pull or receive via push to do review or perform some other action with data they receive? How will data requestors or apps communicate to the network when data is to be made available to third parties or will the TEF prohibit such activity?

MRTC 2.2.1(iv)(a)(1) - Message Delivery by QHIN on its own behalf:

If a QHIN has business sending or pulling data, shouldn't it be authorized through the same procedures as any other Participant and the underlying organization requesting data be identified as such in the push or query?

MRTC 2.2.2 – Permitted and Future Uses of EHI:

Separating data a Participant originates from data elsewhere obtained seems important, so as not to proliferate duplication of data.

MRTC 2.2.4(i) – Individual Access Services Request; Direct Relationship with QHIN:

How will individuals also provide electronic means to reasonably prove they are who they say they are? Including having done so in the case where a third party app is asserting Individual Access Service request on behalf of an individual is also important.

QTF Overview:

As the ONC's proposed rule requires the use of FHIR for Health IT certification, and presumably as the preferred protocol for end user access to data, mandating IHE protocols for QHIN-to-QHIN communication creates an undesirable dichotomy: an inherently document-based inter-QHIN backbone based on IHE protocols, and a non-document-based front-end based on FHIR. Since there is not a 100% mapping between the common IHE payloads and FHIR resources, this would introduce an undesirable and unavoidable component of data corruption if data needs to be converted between the backbone document format and the front end FHIR format. As a result, the safest pattern (with respect to data fidelity) that we think would most likely be adopted would be to stick to the lowest common denominator: documents. This would substantially limit the value of using FHIR on the front end, as all extra-QHIN data would likely be available only in its document form (e.g. as FHIR DocumentReference and Binary resources), even if the data source was in possession of richer data that could have been expressed more completely in the form of FHIR resources.

Additionally, since a QHIN would need to respond to extra-QHIN requests using the document-based IHE protocols, this creates a substantial disincentive to using a non-document-based framework like FHIR within a given QHIN. Since no "perfect" conversion exists between documents and FHIR resources, documents would be the de facto packaging of data for both intra-QHIN and inter-QHIN exchange, since this would be the minimum required.

QTF Query Scenario – SAML assertions:

Using Direct without adaptation, purpose (and any other necessary elements) can be asserted using the Direct Project Context IG. As FHIR transactions are layered on, an analogous assertion can be made

using signed JWTs. Another benefit of using Direct as implemented today is that patients not participating in a QHIN would be capable of receiving messages pushed through a QHIN.

QTF Message Delivery:

It seems there are providers who do not yet use Direct. Will building out this new infrastructure make that happen? If you can't find a Direct address for a provider, they probably don't want to be getting referrals, so it seems kind of pointless to create a parallel network to try to reach them. It's difficult to understand what purpose this new modality would serve. The earlier suggestion of service addresses that can behave as HIE-wide senders to established receivers seems more practical. If you want to force the ability to receive referrals, it seems less expensive to ask those providers to activate the already-built-in capability of using Direct at their endpoints rather than building yet another transport modality to attempt to address the lack of participation. If the issue is usability of Direct that is already in place, similarly it would be easier to address the usability of this technology that is already in place rather than create something entirely new.

Secure Texting probably makes sense to consider as an alternative message delivery standard as well.

ONC Request for Comment #1: Should the QTF specify additional standards or approaches for securing QHIN Exchange Network transactions (e.g. OASIS Web Services Security)?

No. Additionally, while we agree with TLS as a baseline requirement to promote interoperability, it seems that a QHIN should also be permitted to establish an equivalently secure connection with another QHIN if it results in more efficient use of network resources, which could be a substantial consideration at scale.

QTF Mutual QHIN Server Auth:

If this architecture is required, Secure Texting is of particular interest for message delivery since this is an established standard and lends itself well to this modality. Otherwise, the sending QHIN should send a Direct message directly to the known recipient.

ONC Request for Comment #2: What specific elements should a SAML assertion for User Authentication include?

The identifying information would ideally include sufficient information to uniquely identify the Participant, Participant Member, and Individual User, as applicable. Note that not all transactions will involve an individual user. The specific formatting of this data, e.g. SAML/XUA vs JWT vs other, should be left to the RCE as different assertion formats may make sense for different exchange protocols.

ONC Request for Comment #3: Should QHINs be required to transmit other authorization information (e.g., user roles, security labels) in addition to Exchange Purpose and any information required by IHE XUA? What specific elements should a SAML assertion include?

If this is not going to be rolled out until sufficient security tagging has been established, it might make sense to focus that effort on FHIR R4 with assertions in signed JWTs rather than IHE queries with SAML assertions and using a standardized patient match operation supported by additional and verified 2020 Edition patient attributes, if a national patient identifier is not yet established or as we move toward one.

QTF QHIN Exchange Network Message Delivery as complement to Direct Messaging:

There are easier ways to achieve a query in this case—using tools already in place. The QHIN can stand up a Direct service address if the Participant or Participant Member does not have one of their own. The problem where the receiver's address cannot be located is solved with the ONC's requirement of directory listing in the requirements of certified Health IT. If the receiver does not have a Direct address, they are simply not equipped to deal with electronic referrals. Creating a new transfer modality will dilute interoperability by forcing Health IT companies to build more technology to solve the same problem.

QTF Table 8:

Add Secure Texting as an Alternative/Emerging Standard/Profile for Message Delivery.

QTF Message Delivery with XCDR:

There appear to be several limitations to the proposed alternate message delivery protocol. We would expect that most QHINs would desire to enable the message delivery solicitation step using XDR, as is described in the XCDR specification. When used this way, however, XCDR requires the XDR document source to know the target's homeCommunityId, which means that either the protocol must be customized or a directory of all possible recipients (including patients) and their homeCommunityId would need to be constructed, accessible to every first degree entity that is constructing the outgoing XDR message, and new protocols for aggregating and maintaining such directories created. In the case where the first degree entity is actually an HIE, this same data may need to be available downstream if that is where the XDR messages are constructed.

There is no push equivalent of broadcast query defined in the XCDR specification – the initiating gateway needs to know the homeCommunityId of the recipient and which QHIN to send the data to – so a message could not simply be submitted blindly to the network without this information. If the downstream sender cannot resolve the community ID of the recipient(s) on its own using a directory as in the previous paragraph, the alternative would be to design a new "edge protocol" between the first degree entity and the QHIN to communicate minimum required destination metadata as well as a mechanism for a QHIN to resolve that information into a target homeCommunityId.

Additionally, most existing EHR implementations of XDR are synchronous. The delays involved in cross-community exchange may result in substantial overhead as connections are kept open awaiting end-to-end responses. There is no optimal solution to this issue as allowing a longer response time would substantially constrain available network resources while shortening the timeout interval at all would result in unacceptable HTTP request time-out failures, especially if bottlenecks were encountered at any of the intermediate relay paths. It is non-trivial to update synchronous implementations to support asynchronous exchange.

ONC, the Health IT community, and numerous other stakeholders have invested a large amount of time and resources into developing and scaling Direct messaging as a robust "push" technology. Due to its inclusion in the last two major ONC Health IT certification program editions, Direct technology is widely available today and is broadly deployed. Natively asynchronous protocols like Direct can better absorb spikes in traffic, free local resources more quickly, and can better handle multiple recipients.

Many issues related to cross-organization exchange have been solved over several year. Provider directories have increased in scope and are increasingly accessible. A new protocol would undoubtedly experience the same growing pains and the remaining issues relating to payload incompatibilities would not addressed by simply moving the same payloads through an alternate network.

ONC Request for Comment #7: The IHE XCPD profile only requires a minimal set of demographic information (i.e., name and birth date/time). Should QHINs use a broader set of specified patient demographic elements to resolve patient identity? What elements should comprise such a set?

Yes—best practices are needed here and an industry consortium would probably be helpful. Suggestions for additional metadata to require include: home address, gender, and some other identifier like a plan ID, SSN last 4, email address, or telephone number. Ideally unique identifiers like an OpenID identifier will be added to patient records as well. A smaller number of data points may be workable if this information has been verified by one party or the other (before being added to the patient index or before part of a matching query). The FAST Identity Tiger Team is preparing to interview SMEs on this topic and additional information may be found here: https://oncprojectracking.healthit.gov/wiki/display/TechLabSC/Identity+Tiger+Team#identity

ONC Request for Comment #8: There are many possible approaches to Patient Identity Resolution, each with its own benefits and risks. For example, a centralized index of patient identity information may be more efficient for resolving patient identities across disparate communities, but also poses a greater risk to privacy if the system is compromised. Federated approaches may be less susceptible to external threats like cyberattacks, but harder to scale across many communities. Recognizing that new technologies and business entities with robust identity matching solutions may disrupt traditional approaches, should the QTF specify a single standardized approach to Patient Identity Resolution across QHINs?

Probably not since this would not encourage innovation, but recommending best practices in the form of minimum data sets and a rubric for scoring matches is certainly needed.

ONC Request for Comment #9: Different communities tolerate different degrees of risk with respect to accurately matching patient identities. Should QHINs meet a minimum performance standard (e.g., a minimum acceptable matching accuracy rate) over a specified time period? Likewise, different algorithmic techniques for matching patient identities use different approaches and must be tuned to the applicable patient population and continuously refined over time. Should QHINs measure and report on the performance of the algorithm(s) they rely on (e.g., by calculating precision, recall, etc.)?

Ideally yes, but this may vary by use case. An individual patient's query for their own data should be a 100% confidence match (the patient can be required to embed their OpenID identifier into all their records they want to be able to access through a QHIN). A result being returned to a Covered Entity may be a lower confidence level but should include appropriate warnings and informative metadata about the match confidence.

ONC Request for Comment #10: Recognizing there are different ways to implement Record Location services, should the QTF specify a single standardized approach across QHINs?

If it does not, cross-testing should be performed to evaluate performance across systems.