eHealth Exchange™ Document Submission Web Services Specification

Version 3.0

Revision 008

Published: May 19, 2021

Copyright© 2020-2021 All rights reserved.

Contributors

Name	Organization	Role
Jay Nakashima	eHealth Exchange	Executive Sponsor / Contributor
Eric Heflin	Heflin Consultancy LLC	Author
Joe Lamy	AEGIS.Net	Lead Author
John Moehrke	By Light Consulting	Architectural Contributor
Mike Yackanich	eHealth Exchange	Contributor
Mike McCune	eHealth Exchange	Contributor
Matt Blackmon	The Sequoia Project	Contributor
John Donnelly	IntePro Solutions	Contributor
Alexandra E. Voss		Proof reading
Jeff Helman	AEGIS.net	Contributor
Didi Davis	The Sequoia Project	Contributor

Document Change History

Version	Date	Changed By	Items Changed Since Previous Version
3.0 r001	2020-01-19	Eric Heflin	Initial use case, functional requirements, technical specification and architectural diagrams.
3.0 r002	2020-05-13	Joe Lamy	Initial Draft in new format. Use cases content creation. Updated sequence diagram. Moved and edited prior NwHIN Doc Submission text and concepts. Created aux features matrix (to be incorporated).
3.0 r003	2020-05-20	Eric Heflin	Content changes. Added Security considerations, Auditing, Network-network considerations, eHx

Copyright© 2020-2021 All rights reserved.

			Hub considerations, and Sequoia test tool considerations initial text.
3.0 r004	2020-05-28	Joe Lamy	Added example messages, Informative specifications, Roadmap/Backlog, Operational Considerations. replaced HITSP C80 with links to value sets maintained in FHIR® R4. Refined many other sections.
3.0 r005	2020-08-02	Joe Lamy	Completed draft text for initial scope as well as Persistence and Provenance. Added requirements for compatibility with the TEFCA QTF. Added potential Hub translations between XCDR and XDR transactions. Added Push Comparison Table.
3.0 r006	2021-01 to 04	Joe Lamy	Added FHIR mechanisms.
3.0 r007	2021-04-27	Eric Heflin	Final edits prior to publishing draft for broad review by the full eHealth Exchange community.
3.0 r008	2021-05-13	Joe Lamy/Eric Heflin/Ali Voss	Final edits to final version being submitted to the eHealth Exchange Coordinating Committee for formal approval as a new DURSA-recognized performance-and-service specification.

Document Approval

Version	Date	Approved By	Role
3.0 r007	2021-05-07	eHealth Exchange Technical Work Group	Approves all specifications for production use
3.0 r008	2021-05-18	eHealth Exchange Coordinating Committee	Approves all specifications for production use

Copyright© 2020-2021 All rights reserved.

Table of Contents

Copyright© 2020-2021 All rights reserved.

Copyright© 2020-2021 All rights reserved.

Copyright© 2020-2021 All rights reserved.

Copyright© 2020-2021 All rights reserved.

Copyright© 2020-2021 All rights reserved.

1 PREFACE

1.1 Introduction

The eHealth Exchange (eHx) Web Service Interface specifications define the core set of standard services to be implemented by each Participant on the eHx network in order to securely exchange interoperable health information over the public Internet. The eHx consists of Health Information Exchanges, federal, state and county agencies, large Integrated Delivery Networks, Dialysis clinics, pharmacies, practices and more. The eHx also connects to other networks across the United States.

The eHx functional services provide discovery and information exchange capabilities and rest upon a foundational set of messaging, security and privacy services.

This document presents the eHx Document Submission Web Service Interface specification version 3. The purpose of this service is to allow one Participant on the eHx to securely "push" data for a given patient to another Participant in a manner designed for automated process of interoperable clinical content.

1.2 Intended Audience

The primary audiences for this eHx specification are individuals responsible for implementing software solutions such as project managers, CTOs, CISOs, software engineers, technical product managers, IT managers, operations staff, and others with similar roles.

It's assumed that the reader has a moderate degree of familiarity with IHE profiles, esp. XCPD, XCDR, XDR, XCA, XDS, ATNA, CT and XUA. The reader should also have awareness of foundational OASIS standards such as WS-Addressing and SAML.

1.3 Typographical Conventions

Quoted items are used precisely in this document, following computer science standards, and only quote exact text. So, for example, a quote that is followed by punctuation, is quoted as:

"https://example.org/csrk/index.html". not as "http://example.org/csrk/index.html."

Copyright© 2020-2021 All rights reserved.

This convention is followed since the punctuation often changes the semantic meaning of the quoted text and can introduce an error, as it would in the above example.

1.4 Focus of this Document

This document defines the eHx Document Submission Web Service Interface Specification. The purpose of this service is to provide the ability to "push" data for a given patient from one Participant to another. This "directed push" is a different model of exchange than subscription (also called "push notifications") because, with directed push, the sender explicitly sends the data to a receiver, while, with subscriptions, the receiver is just one of potentially many subscribers. Another potential deployment is that directed push could implement a subscription where the subscription mechanism is managed out-of-band.

1.5 Business Needs Supported by this Specification

This specification is intended to create an implementation-quality technical specification designed to meet the following business requirements:

1. Security
2. Enabling automation, audit logging and access control
3. Enabling various eHealth Exchange use cases whereby which one Participant needs to convey information to a second Participant via an outbound "push" from the first Participant
4. Allowing sending Participants to include metadata with the transaction
5. Allowing sending Participants to include payloads with the transactions, such as a FHIR bundle or a C-CDA XML document
6. Providing for use cases where there is no need for technical pre-coordination (such as electronic subscription management) before the sending Participant transmits the data

1.6 Scope

This specification defines:

1. The ability of one Participant to push data for a patient to another Participant using SOAP web services or HL7 FHIR mechanisms for transport, and HL7 CDA documents, FHIR documents, raw documents (e.g. PDF) or FHIR resources for content,
2. The ability for a patient to push data, as long as it is through a sending intermediary that is a participant on the eHealth Exchange,
3. Conditional persistence and provenance requirements for pushed data,

Copyright© 2020-2021 All rights reserved.

4. Compatibility requirements for interacting with participants on the TEFCA QTF, and
5. eHx Hub behavior to perform translations between different push transactions.

This specification does not define:

1. The ability to push data for multiple patients,
2. The ability to push data for a deidentified patient,
3. The ability to push aggregate data not tied to specific patients,
4. The ability to push data unrelated to patients,
5. The ability to push data in response to a subscription,
6. The ability to push data using the Direct protocol (this is specified elsewhere),
7. The ability to push data using FHIR Messaging (however a placeholder is provided for future capabilities),
8. The ability to search data using POST,
9. The ability for a patient to push data directly to a participant on the eHealth Exchange (e.g. via a smartphone app or consumer device), and
10. eHx Hub behavior to perform these transactions across multiple Participants for the same document (e.g., broadcast delivery).

Note: Provenance information is used to track authorship and identification of entities that have contributed to or changed data. In a Push context, provenance responsibilities can fall on both the Initiating and Receiving Participants. The requirements here do not fully encompass the provenance responsibilities of a system; they only cover the cases of pushing information to or receiving information from an external system. A comprehensive Provenance specification may eventually be written.

1.7 Outstanding Issues

The authors are working with outside standards bodies and exchange workgroups on a number of outstanding issues that impact this specification. Our goal is to limit the impact these “in-flight” topics have on participants, so in writing this specification:

- We have highlighted areas of concern and instability in the standards.
- We have attempted to capture current understanding, and have provided guidance that we believe is likely to hold up over time.
- We have limited the coverage of certain topics so that we may provide full analysis and guidance in future specifications and policy.

The following are the key areas of concern.

Copyright© 2020-2021 All rights reserved.

1.7.1 Federation and Routing

In SOAP web services used for healthcare, while there are many competing/complementary mechanisms to accomplish federation and routing, for the most part usage has coalesced around IHE-constrained SOAP and WS-Addressing, the IHE concept of a “home community id” for federated deployments, and eHx/IHE-constrained SAML to enable access control for federated requesters.

As participants deploy ever more complex subnetworks and exchanges begin to bridge together, they are finding that use cases and goals they had but had never made explicit are not being met due to limitations in standards or operational configuration. Working with outside groups, we are helping to finally make these needs explicit and update guidance accordingly.

During the writing of this specification we have worked with IHE to define a consistent way to push to federated systems, whether using XDR, XCDR, MHD, Direct, or a combination. See ITI MHD to a Federation work item (https://github.com/IHE/IT-Infrastructure/issues/142). This work is incorporated in this specification.

We are also working with HL7 on the new UDAP FHIR IG, with IHE on updates to the IUA profile, and with Carequality on scenario identification, to support access control use cases when there are federated initiators. This work will be incorporated in the separate eHx Authorization Framework specification.

1.7.2 Provenance

Provenance is not currently well understood or implemented widely. There are disconnects between what is possible, what producers want to provide, and what consumers want to receive.

We have made use of the recently written HL7 Basic Provenance guide, and have been working with relevant experts on adding additional capabilities introduced in the HL7 Data Provenance guide and the IHE Reconciliation profile. While this work is in this specification, it is still in flux and all or some of it will be moved to a separate specification.

1.7.3 FHIR resource life cycles

We have identified the need for guidance on the life cycle of FHIR resources. This will likely be added in a separate specification.

1.8 Content Type Policy

The content type policy goals are as follows, to:

Copyright© 2020-2021 All rights reserved.

1. Ensure new content is **generated** in the most preferred format possible.
2. Generate and exchange the use of structured formats when possible.
3. Advocate for most contemporary standard being the most preferred approach pervasively throughout the eHx.

In the absence of policy to the contrary, the eHealth Exchange requires that newly created content to be exchanged MUST be generated using structured formats when such an industry standard exists. When no such format exists, then, and only then, Participants can use semi- and unstructured formats such as spreadsheets, csv files, PDFs, raw XML and raw JSON.

For some use cases, eHx specifications point to a single defined content format for a given use case, e.g., Patient Corrections that use a Task resource.

In the event that multiple standard structured formats could be used to accomplish the same use case, then Participants MUST generate **new** content using the following precedence, highest listed first:

1. FHIR resources or bundles of resources (including FHIR documents) in XML format
2. FHIR resources or bundles of resources (including FHIR documents) in JSON format
3. C-CDA 2.1 documents
4. C-CDA 1.1 documents
5. HITSP C32
6. HL7 v2.x

Participants SHOULD exchange **legacy** content using these preferred formats, but any applicable format may be used. Example: C32s had already been generated and live in a repository. Example: legacy paper documents could be captured as PDFs or wrapped in a CDA (e.g. C62) with an unstructured body (both with corresponding XDS metadata). This is in recognition that it is better to exchange content in any format than not to exchange it at all.

Participants will negotiate formats as follows:

- For the Pull mechanism, Initiating Participants SHOULD support receiving content in the more preferred formats above and MUST use the MIME type and format code returned in the Query for Documents response to enforce this precedence.
- For the Push mechanism,
    - Receiving Participants SHOULD support receiving content in the more preferred formats above and MUST ensure that supported formats are expressed in the eHx directory per current policy and capabilities.

Copyright© 2020-2021 All rights reserved.

- Initiating Participants SHOULD use the eHx directory to identify the highest precedence format supported by the Receiving Participant, and if identified, MUST push in that format.

1.9 Content Equivalency Policy

CONF-001: If the Receiving Participant provides multiple APIs (e.g. FHIR, SOAP) to the same underlying data, they SHALL provide substantially similar data irrespective of the API used.

Informative: For example, if a system accepts and persists a document through the SOAP web services API, and that system makes documents available through SOAP and FHIR APIs, it shall make that same document available via both APIs, although the format may be different.

1.10 Intellectual Property Rights

This document is copyright© 2020-21 eHealth Exchange.

HL7®, CDA®, and FHIR® are registered trademarks of Health Level 7, and are used with permission.

1.11 Related Documents

1.11.1 Normative

This interface specification references the following standards for normative requirements. Specific deviations from or constraints upon these standards are identified below.

11. IHE IT Infrastructure Technical Framework
    1. Org/SDO: Integrating the Healthcare Enterprise (IHE)
    2. Version #: Revision 17.0 (2020-12-02)
    3. Links to documents and sections of focus:
        1. **[IHE ITI TF-1]** Volume 1: https://profiles.ihe.net/ITI/TF/Volume1/index.html
            1. Section 2.1 – Dependencies among Integration Profiles
            2. Section 2.2.15 – Cross-Enterprise Document Reliable Interchange (XDR) Brief description
            3. Section 15 – XDR Full documentation and use cases
        2. **[IHE ITI TF-2]** Volume 2: https://profiles.ihe.net/ITI/TF/Volume2/index.html
            1. Section 3.41 – Provide and Register Document Set-b (ITI-41)

Copyright© 2020-2021 All rights reserved.

2. Section 3.80 – Cross-Gateway Document Provide (ITI-80)
3. Appendix V – Web Services for IHE Transactions
4. Appendix Z – FHIR Implementation Material.

3. **[IHE ITI TF-3]** Volume 3: https://profiles.ihe.net/ITI/TF/Volume3/index.html
    1. Section 4 – Metadata used in Document Sharing profiles
4. eHx Deviations or Constraints:
    1. Only certain types of metadata are normatively adopted. See SOAP Web Services Push (TBD link) API Description for details.
5. Underlying Specifications:
    1. ebXML 3.0: consists of
        1. ebRS 3.0: http://docs.oasis-open.org/regrep/regrep-rs/v3.0/regrep-rs-3.0-os.pdf
        2. ebRIM 3.0: http://docs.oasis-open.org/regrep/v3.0/specs/regrep-rim-3.0-os.pdf
    2. MTOM 25 January 2005: http://www.w3.org/TR/soap12-mtom/
    3. XOP 25 January 2005: http://www.w3.org/TR/2005/REC-xop10-20050125/

12. IHE IT Infrastructure Technical Framework Supplement: Cross-Community Document Reliable Interchange (XCDR)
    1. Org/SDO: Integrating the Healthcare Enterprise (IHE)
    2. Version #: Revision 1.4 – Trial Implementation (2017-07-21)
    3. Links to documents and sections of focus:
        1. IHE XCDR Supplement https://www.ihe.net/uploadedFiles/Documents/ITI/IHE_ITI_Suppl_XCDR_Rev1.4_TI_2017-07-21.pdf
        2. Note: we do not refer to this supplement directly; it describes content to be integrated in the existing ITI Volumes. When we refer to that content in this specification, we use the destination volume, e.g., **[IHE ITI TF-2c]**.
    4. eHx Deviations or Constraints:
        1. Only certain types of metadata are normatively adopted. See SOAP Web Services Push (TBD link) API Description for details.
        2. Only ITI-80 from this supplement is supported. The XDR Transmit Home Community ID option on ITI-41 is not supported.
    5. Underlying Specifications: None
13. IHE IT Infrastructure Technical Framework Supplement: Mobile access to Health Documents (MHD) With XDS on FHIR
    1. Org/SDO: Integrating the Healthcare Enterprise (IHE)
    2. Version #: TBD

Copyright© 2020-2021 All rights reserved.

1. The most recent published version is 4.0.0-comment – the Public Comment release for the first IG authored version – but this specification is being written towards the next version, which will incorporate numerous changes, many of which were prompted by this specification. The link to the MHD FHIR IG below is to the current build. A future version of this specification may point to the resulting normative version.

3. Links to documents and sections of focus:
    1. **[IHE MHD IG]** Current build of MHD FHIR IG: http://build.fhir.org/ig/IHE/ITI.MHD.
4. eHx Deviations or Constraints: See FHIR Push API Description for details.
5. Underlying Specifications: [HL7 FHIR R4]

14. IHE IT Infrastructure Technical Framework Supplement: IHE Appendix on HL7® FHIR®
    1. Org/SDO: Integrating the Healthcare Enterprise (IHE)
    2. Version #: Revision 2.2 – Trial Implementation (2020-08-28)
    3. Links to documents and sections of focus:
        1. IHE Appendix Z Supplement https://www.ihe.net/uploadedFiles/Documents/ITI/IHE_ITI_Suppl_Appx-Z.pdf
        2. Note: we do not refer to this supplement directly; it describes content to be integrated in the existing ITI Volumes. When we refer to that content in this specification, we use the destination volume, e.g., **[IHE ITI TF-2x]**.
    4. eHx Deviations or Constraints: None
    5. Underlying Specifications: [HL7 FHIR R4]
15. **[HL7 FHIR R4]** Fast Healthcare Interoperability Resources (FHIR®)
    1. Org/SDO: Health Level 7 (HL7®)
    2. Version #: R4 (4.0.1): http://hl7.org/fhir/R4/
    3. Links to pages of focus:
        1. Value set binding strengths
        2. classCode value set
        3. confidentialityCode value set
        4. eventCodeList value set
        5. formatCode value set
        6. healthcareFacilityTypeCode value set
        7. practiceSettingCode value set
        8. typeCode value set
    4. eHx Deviations or Constraints:
        1. Value sets are further constrained. See SOAP Web Services Push API Description and FHIR Push API Description for details.

Copyright© 2020-2021 All rights reserved.

5. Underlying Specifications: These value sets were previously maintained at http://www.hitsp.org/ConstructSet_Details.aspx?&PrefixAlpha=4&PrefixNumeric=80 by HITSP, which is no longer an active standards body.

16. **[HL7 Basic Provenance]** HL7® Guidance: Basic Provenance for C-CDA and FHIR®
    1. Org/SDO: Health Level 7 (HL7®)
    2. Version #: Release 1 - US Realm: http://www.hl7.org/implement/standards/product_brief.cfm?product_id=531
    3. eHx Deviations or Constraints:
        1. See Provenance sections for details.
    4. Underlying Specifications: [HL7® Data Provenance]
17. **[HL7 Data Provenance]** HL7® CDA® R2 Implementation Guide: Data Provenance
    1. Org/SDO: Health Level 7 (HL7®)
    2. Version #: Release 1 - US Realm: http://www.hl7.org/implement/standards/product_brief.cfm?product_id=420
    3. eHx Deviations or Constraints:
        1. See Provenance sections for details.
    4. Underlying Specifications: None
18. **[IHE RECON]** IHE Patient Care Coordination Technical Framework Supplement: Reconciliation of Clinical Content and Care Providers (RECON)
    1. Org/SDO: Integrating the Healthcare Enterprise (IHE)
    2. Version #: Rev. 3.2 – Trial Implementation: https://www.ihe.net/uploadedFiles/Documents/PCC/IHE_PCC_Suppl_RECON_Rev3.2_TI_2016-11-11.pdf
    3. eHx Deviations or Constraints:
        1. See Provenance sections for details.
    4. Underlying Specifications: None
19. Trusted Exchange Framework and Common Agreement (TEFCA)
    1. Org/SDO: The Office of the National Coordinator for Health Information Technology (ONC)
    2. Version #: Draft 2 (2019-04-09): https://www.healthit.gov/sites/default/files/page/2019-04/FINALTEFCAQTF41719508version.pdf
    3. Sections of focus:
        1. **[TEFCA QTF]** Appendix 3: Qualified Health Information Network (QHIN) Technical Framework, Draft 1

1.11.2 Informative

This interface specification references the following standards for informative guidance:

Copyright© 2020-2021 All rights reserved.

20. **[HL7 eICR]** HL7® CDA® R2 Implementation Guide: Public Health Case Report, Release 2 - US Realm - the Electronic Initial Case Report (eICR) - referenced by high level use cases
    1. http://www.hl7.org/implement/standards/product_brief.cfm?product_id=436
21. **[HL7 RR]** HL7® CDA® R2 Implementation Guide: Reportability Response, Release 1, STU Release 1.0 - US Realm - referenced by high level use cases
    1. http://www.hl7.org/implement/standards/product_brief.cfm?product_id=470

1.12 Relationship to Other eHx Specifications

This specification is related to other eHx specifications as described below.

The **Messaging Platform** specification specifies a base set of messaging standards and web service protocols that must be implemented by each eHx gateway and applies to all transactions. All eHx inter-nodal messages are SOAP messages over HTTP using web services and must be encrypted and digitally signed.

Informative: The eHx expects to adopt HL7® FHIR®-based transactions in the future.

The **Authorization Framework** specification defines the exchange of metadata used to characterize each eHx request. The purpose of that exchange is to provide the responder with the information needed to make an authorization decision for the requested function and to provide high-resolution audit-logging. Each initiating message must convey information regarding end user attributes and authentication using SAML 2.0 assertions.

Together, the Messaging Platform and the Authorization Framework define the foundational messaging, security and privacy mechanisms for the eHx.

The **eHx Directory Specification** defines the API for using the eHx Directory. In this transaction, Initiating Participants look up Receiving Participants in the directory.

The **eHx Hub Specification** defines common behavior of the eHx Hub across all transactions. Hub behavior specific to this transaction is included in this document and referenced by the Hub Specification.

In all cases, the data exchanged between Participants will involve the communication of individually identifiable health information (defined in 45 CFR Parts 160, 162 and 164). When individually identifiable information is exchanged, then each Participant must have a common understanding of the patient's identity. To facilitate a common understanding and prepare two Participants for spontaneous exchange of data that identifies a patient, the Participants may utilize the eHx **Patient Discovery** Interface specification to share the identity of a patient

Copyright© 2020-2021 All rights reserved.

between the exchanging Participants (see Section 3.9.4 for further details on the use of a patient identifier).

1.13 Conventions Used in this Specification

The keywords "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119].

Conformance statements (**CONF-XXX:...**) are normative requirements. Each conformance statement is intended to be a context-independent, testable assertion that can be directly tested by a human or an automated test case.

Informative blocks of text (Informative:...) are explanatory and not binding in any way.

2 High Level Use Case Descriptions

2.1 Use Case Purpose

The focus in these use cases is to identify realistic and needed scenarios that drive specific feature needs from the Push capability, e.g., routing, patient identity and exception handling. Use cases are a critical component of the specification as they are our benchmark; with use cases we can determine if the associated functional requirements and the technical implementation specification are sufficient to meet our shared requirements. They also allow us to determine if there are gaps or unneeded features. Hence properly defined use cases are critical to the success of this implementation guide.

2.2 Use Case: Public Health Reporting Through Intermediary

A multi-state association of hospitals provides reportable conditions to several state recipients using Electronic Case Reporting (eCR).

A patient presents at a provider with symptoms consistent with COVID-19. The provider performs tests and records a clinical diagnosis in the patient's record. The EHR system evaluates the diagnosis against a set of trigger codes for reportable conditions and finds a match. The EHR system creates and pushes an Electronic Initial Case Report (eICR) [HL7 eICR] document for the patient to a public health reporting service, which acts as an intermediary to state Public Health

Copyright© 2020-2021 All rights reserved.

Agencies (PHAs). The reporting service is a single clearinghouse already known to the EHR system, so no directory lookup is needed at the time of the event.

The reporting service opens the eICR and, based on the state rules for reporting, inspects the codes, patient address and facility location, and determines the reportability of the condition, as well as the jurisdiction (i.e. which PHAs) to route it to. The reporting service pushes the eICR accordingly. If needed, the reporting service looks up PHAs in a service directory.

The reporting service then creates an HL7 Reportability Response (RR) CDA document [HL7 RR] and pushes it to the appropriate PHAs as well as back to the original EHR system. For this step, the reporting service needs to have retained the identity of the original sending system, which it uses to look that system up in a service directory.

2.2.1 Assumptions

There is no assumption that either the reporting service or any PHAs previously knew the patient. Thus, the sending EHR does not need to match to a known patient at either the reporting service or the agencies. The agencies will identify the patient locally if necessary based on demographics in the eICR. For example, this could be done to prevent duplicate reporting for the same patient event.

The sending EHR does not need to know about routing beyond the reporting service.

The reporting service looks up the “return address” for the RR using the Home Community ID of the sender in the metadata of the Push transaction.

The eICR and RR documents do not need to include full metadata when they are pushed.

The eICR and RR documents are not assumed at any of the destinations to be stored and made available for future queries.

Copyright© 2020-2021 All rights reserved.

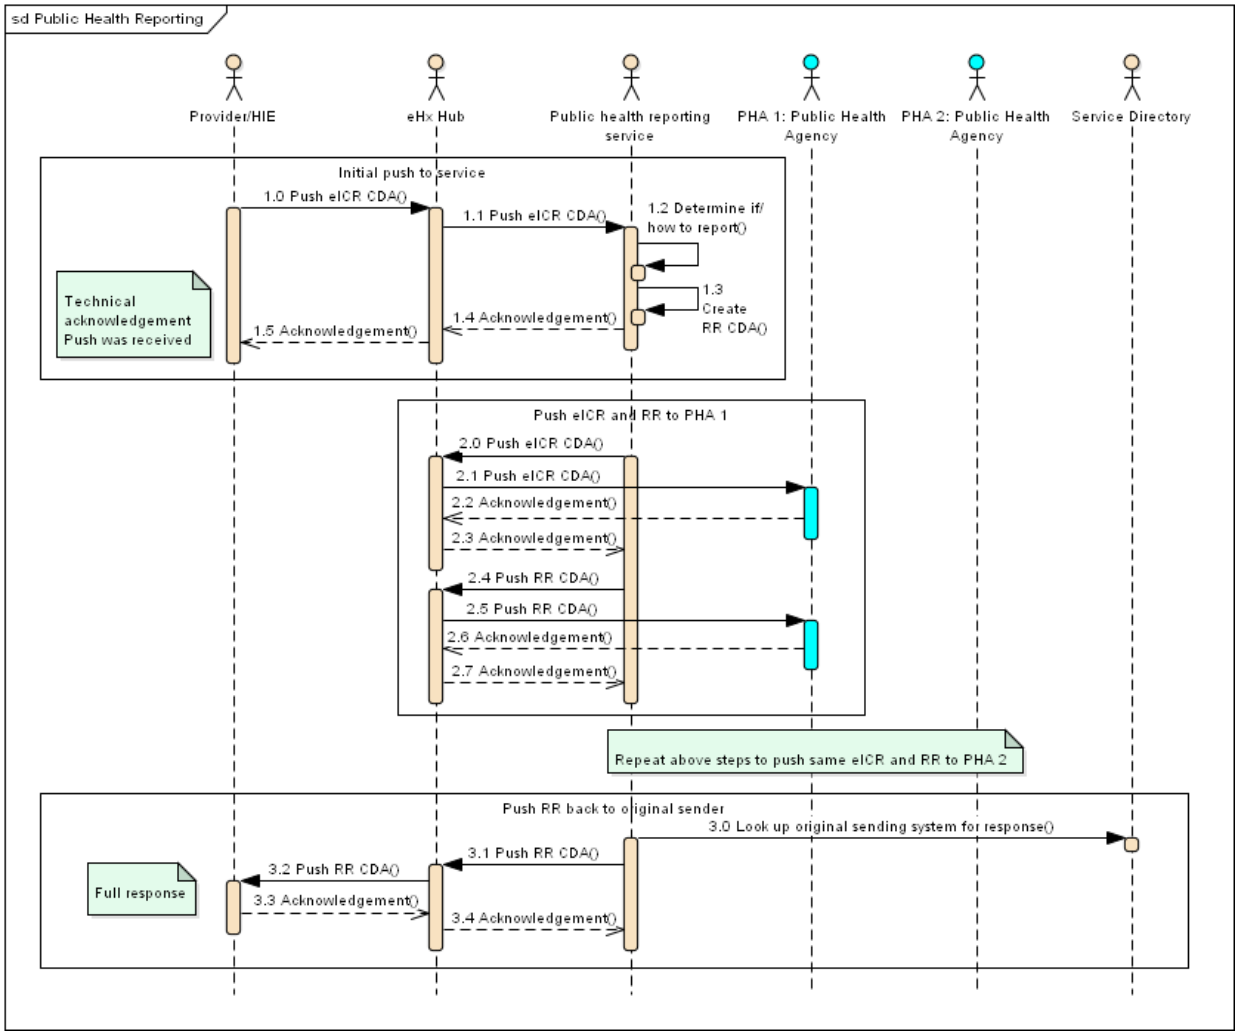

Figure 1 Public Health Reporting Through Intermediary

2.3 Use Case: Reporting Clinical Events to Patient Home Community

There are a number of scenarios whereby a patient obtains care outside of their normal patient home community, and those clinical events should be reported back to the patient home community to be incorporated into the patient's record. For example:

1. A diagnosis of a communicable condition such as COVID-19 is confirmed.
2. A nationwide pharmacy has an agreement with a federal agency to report immunizations for that agency's patients to the agency.

For this use case we will detail the immunization reporting example.

Copyright© 2020-2021 All rights reserved.

Precondition: the nationwide pharmacy has an agreement with a federal agency to report to that agency any immunizations for that agency's patients.

A patient gets an immunization at a pharmacy. The patient presents an identification card with their ID as known by the agency. The pharmacy system attempts to match the patient at the agency, passing the agency's patient ID as well as other patient demographics to the agency. Note that use of this ID greatly increases the chances of a positive match.

Finding the patient, the pharmacy system creates a document containing the encounter and immunization and pushes it to the agency. The immunization encounter document may include full metadata.

The agency has sufficient information to persist this document in the patient's record and make it available for future queries. Alternatively, the agency may import the clinical content into the patient's record and make it available through natively generated documents.

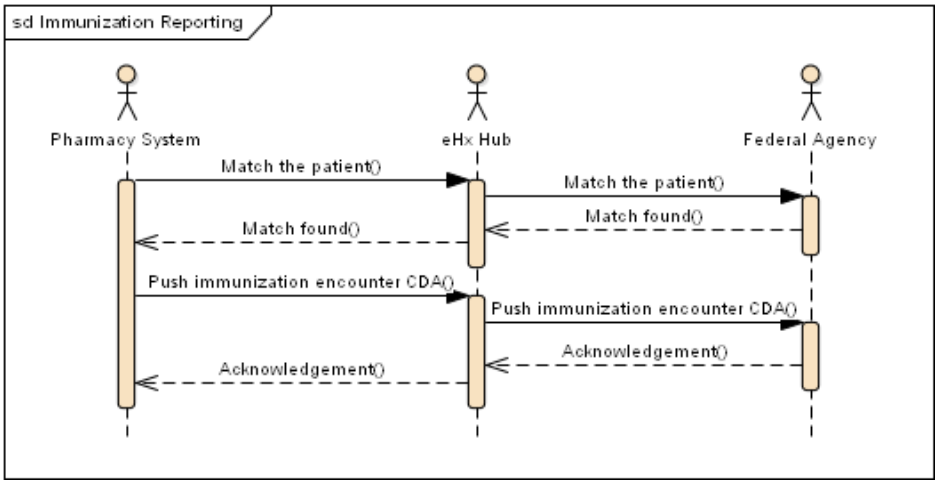

Figure 2 Reporting Immunization Events to Patient Home Community

3 SOAP Web Services Push API Description

3.1 Changes Since Document Submission 2.0

CONF-002: The Deferred mode for Document Submission is deprecated; it SHALL NOT be used.

Informative: prior versions of this specification defined a Deferred mode, but it was not based on an equivalent mechanism in the underlying IHE transaction and has been deprecated.

Copyright© 2020-2021 All rights reserved.

CONF-003: The use of Document Submission to push De-identified Documents is deprecated; it SHALL NOT be used.

Informative: prior versions of this specification defined support for this, but the mechanism would require additional details to be implementable, and a use case has not been presented.

Informative: The remaining changes in the updated specification are summarized as follows:

- Adopted latest IHE specifications
  - Added XDR Metadata-Limited Document Source
- Added XCDR mechanism
- Added high-level business use cases and system-level use cases
  - Included use of eHx Directory in workflow
- Added provenance requirements
- Defined value sets for coded metadata fields in submissions
- Clarified prior patient matching
- Clarified how to obtain references to existing metadata
- Clarified submitting updates to documents, including reflecting updates in metadata as well as CDA content
- Defined allowed/disallowed variations in submission (e.g., folders)
- Addressed special case of cross-community cross-author updates
- Clarified processing model
- Clarified vague or misleading IHE requirements
- Defined allowed variations in whether/how receiver persists received content and/or makes available for future sharing
- Clarified/expanded error handling
- Clarified conflict detection and resolution
- Added eHx Hub behavior
- Analyzed and ensured compatibility with the TEFCA QTF
- Added Directory requirements
- Added new examples with full feature coverage
- Added feature comparison table, comparing various flavors and versions of push

3.2 Definitions

In this interface specification, a “document” refers to the format of clinical data as it is transferred between Participants and not as it is stored within a Participant system or specific electronic health record (EHR) system. A Participant and its participating organizations may store clinical data in whatever format or repository they choose. Specifically, a “document”

Copyright© 2020-2021 All rights reserved.

transferred between Participants need not meet the criteria for persistence, stewardship, etc., as identified by the HL7 Structured Documents Working Group.

The following terms are defined based on which transaction option of the interface (XDR or XCDR) is implemented:

- The "Document Submission transaction" is a "push" of documents and metadata from an Initiating Participant to a Receiving Participant.
  - **CONF-004**: For the XCDR option, Document Submission SHALL correspond to the IHE ITI-80 Cross-Gateway Document Provide transaction.
  - **CONF-005**: For the XDR option, Document Submission SHALL correspond to the IHE ITI-41 Provide and Register Document Set-b transaction.
- An "Initiating Participant" initiates a Document Submission transaction for one or more available documents on a particular patient.
  - **CONF-006**: For the XCDR option, the Initiating Participant SHALL correspond to the Initiating Gateway actor.
    - Informative: Note that the XDR Metadata-Limited actor is not available for XCDR. Full metadata is required.
  - **CONF-007**: For the XDR option, the Initiating Participant SHALL correspond to either the Document Source actor or the Metadata-Limited Document Source actor.
- A "Receiving Participant" receives a Document Submission transaction.
  - **CONF-008**: For the XCDR option, the Receiving Participant SHALL correspond to the Responding Gateway actor.
  - **CONF-009**: For the XDR option, the Receiving Participant SHALL correspond to the Document Recipient actor.

Where requirements or guidance are identical regardless of the transaction option chosen, the more general Participant term is used. Where requirements or guidance are specific to the transaction option chosen, the IHE profile actor name is used.

CONF-010: In the context of the eHx Patient Discovery Interface specification, the Initiating Participant SHALL correspond to the Initiating NHIO, and the Receiving Participant SHALL correspond to the Responding NHIO.

3.3 Assumptions

The following assumptions underlie this interface specification:

Copyright© 2020-2021 All rights reserved.

- The primary expected use in the context of the eHx is that documents are formatted as XML data following the HL7® Clinical Document Architecture (CDA®) standard (used with permission), but nothing precludes this interface from being used to submit other kinds of documents, such as Adobe Portable Document Format (PDF) files or images.
- The patient to whom the document(s) pertain:
    - Is registered at one or more facilities at the Initiating Participant.
    - Has provided consent to share his or her clinical data, or such consent is not required by the business case under which the Document Submission is occurring; if consent is needed, the mechanism for providing this consent is the subject of the Access Consent Policies specification document.
- How a Participant determines to which other Participant to direct the transaction is not specified.

3.4 Triggers

The Initiating Participant, based on a human decision or an automated workflow, wants to submit document-related information about a patient to a Receiving Participant.

3.5 Transaction Standard

CONF-011: The eHx Document Submission transaction is defined with two transaction options:

- **XDR Option**: This utilizes the IHE ITI-41 Provide and Register Document Set-b transaction for the Cross-Enterprise Document Reliable Interchange (XDR) profile, defined in [IHE ITI TF-2b] 3.41.
- **XCDR Option**: This utilizes the IHE ITI-80 Cross-Gateway Document Provide transaction for the Cross-Community Document Reliable Interchange (XCDR) profile, defined in [IHE ITI TF-2c] 3.80.

The locations and versions of these specifications, as well as other foundational standards for this transaction, are listed in Section 1.11, “Related Documents”.

CONF-012: A Participant MAY support any combination of transaction options in either direction. Informative: For example, the XCDR option as an Initiating Participant and both options as a Receiving Participant.

CONF-013: If a Participant supports both transaction options, it SHALL ensure each function identically except where the transactions inherently differ. Informative: For example, any persistence or error handling implemented for one option must be the same on the other option.

Copyright© 2020-2021 All rights reserved.

CONF-014: Within the transaction options, Participants MAY support the following IHE profile options:

IHE Actor	Supported Options
XDR Document Source	
XDR Metadata-Limited Document Source	
XDR Document Recipient	Accepts Limited Metadata
XCDR Initiating Gateway	
XCDR Responding Gateway	

CONF-015: Unless otherwise specified, Participants SHALL follow all requirements for the respective IHE actors.

Informative: This specification does not include in its scope any grouping of the IHE actors except for the required groupings with ATNA and CT. Participants MAY adopt additional IHE groupings, but no expected behaviors are required.

Informative: There is one available flavor of asynchronous messaging available for ITI-41 and ITI-80, and that is the AS4 Asynchronous Web Services Exchange option. This specification does not support this option at this time; it SHALL NOT be used.

3.5.1 QTF Interoperability

Informative: The [TEFCA QTF] adopts XCDR without conditions, which supports AS4 as an option on both sides. This is not anticipated to be a problem at this time, as it is not required on either side.

3.6 Additional Options

This interface specification defines the following additional options. See the Operational Considerations section of this document for Directory considerations.

- Receiving Participant
  - **Patient ID Required Option**: A Receiving Participant that declares this option indicates that it requires a patient identifier known to it to be included in document metadata. This is not needed in the XDR transaction option unless the

Copyright© 2020-2021 All rights reserved.

XDR Accepts Limited Metadata option is declared, as patient ID is included with full metadata.

- **Persistence Option**: A Receiving Participant that declares this option indicates that it has the ability to persist the documents and metadata that are pushed to it and make them available for subsequent query and retrieval using the eHx Query for Documents and Retrieve Documents transactions. Whether it actually does persist a given document or metadata object is not specified, as this could be subject to security and privacy considerations (e.g. one author pushes a correction to a document from another author, or a patient submits a document).

3.7 System-Level Use Case

3.7.1 Technical Pre-conditions

The following technical pre-conditions exist for this interface specification:

- The document(s) being transmitted pertain to a specific, single patient.

3.7.2 Use Case Steps – "Nominal Flow"

Note: While the entire workflow is described here, the usage of the eHx Directory and the eHx Patient Discovery transaction are not detailed in this specification.

1. This use case begins when the Initiating Participant looks up another Participant that it wishes to push documents to in the eHx Directory.
2. The Initiating Participant obtains the Participant's endpoints for eHx Patient Discovery and a Document Submission transaction option it supports from the Directory. The Initiating Participant examines the options declared by the Receiving Participant on the Document Submission transaction.
3. The Initiating Participant sends a Patient Discovery request to the Receiving Participant to attempt to match the patient by demographics.
4. The Receiving Participant compares the demographics to its known patients and returns a Patient Discovery response to the Initiating Participant. The response contains a single patient match, including demographics and patient ID as known by the Receiving Participant.
5. The Initiating Participant sends a Document Submission request to the Receiving Participant. The request MAY include document entries (i.e. the objects containing document metadata) and the corresponding documents. See ITI-41 and ITI-80 as well as Section 3.9 for metadata requirements.

Copyright© 2020-2021 All rights reserved.

6. The Receiving Participant receives the metadata and the associated document(s), processes them as appropriate, and returns a Document Submission response. See ITI-41 and ITI-80 as well as Sections 3.15 through 3.23 for processing requirements.

3.7.3 Alternate Flows

3.7.3.1 Push to implicitly federated recipient

a. In step 2, the Receiving Participant found in the directory has a Document Submission endpoint that resolves to a federated façade or intermediary Participant, which is in the directory.
b. The use case resumes at step 2, with the following changes:
c. In steps 3 and 5 there is no difference from the Initiating Participant's perspective,
d. In steps 4 and 6, the intermediary Participant forwards the requests to and receives the responses from the actual Receiving Participant, using unspecified means.

3.7.3.2 Push to explicitly federated recipient

1. In step 2, the Receiving Participant found in the directory does not have any endpoints for Document Submission; however, the Initiating Participant locates a "parent Participant" in the directory that supports Document Submission.
2. The use case resumes at step 2, with the following changes:
3. The Parent Participant replaces the original Receiving Participant as the system the Initiating Participant interacts directly with.
4. In step 5, the Initiating Participant includes the Home Community ID for the target (i.e., child) Receiving Participant in the Document Submission request. See section 3.13 Message Requirements: Routing.
5. In step 6, the Parent Participant routes the request to the target Receiving Participant using unspecified mechanisms.

3.7.3.3 Patient ID already obtained

1. In step 1, the Initiating Participant already has the patient identifier.
2. In step 2, the Initiating Participant skips obtaining the Patient Discovery endpoint.
3. The use case resumes at step 5.

3.7.3.4 Patient ID is optional

1. In step 1, a patient ID is optional (see Section 3.9.4, XDSDocumentEntry.patientId).
2. The Initiating Participant chooses to execute one of the following subflows:
    a. Skip the patient match: in step 2, the Initiating Participant skips obtaining the Patient Discovery endpoint. The use case resumes at step 5.

Copyright© 2020-2021 All rights reserved.

b. Attempt the patient match anyway. The use case resumes. If the match fails for any reason, the Initiating Participant MAY end the use case or resume at step 5.

3.7.3.5 No required patient match found

1. In step 4, the Receiving Participant returns no match found, and the patient identifier is required (see Section 3.9.4, XDSDocumentEntry.patientId).
2. The Initiating Participant may attempt to obtain the patient identifier by manual means.
3. If the patient identifier cannot be obtained, the Initiating Participant chooses to execute one of the following subflows:
    a. Resume the use case at step 5, passing the value for sourcePatientId in the patientId attribute.
    b. End the use case.

3.7.3.6 Patient match returned from different community

1. In step 4, the Receiving Participant returns a patient match from a different community.
2. The Initiating Participant looks up the Participant for that community in the eHx Directory, and obtains that Participant's endpoint for a Document Submission transaction option it supports. The Initiating Participant examines the options declared on the Document Submission transaction.
3. The use case resumes at step 5, with the following changes:
    a. The new Receiving Participant replaces the original Receiving Participant.

3.7.3.7 Multiple patient matches returned

Informative: In eHx, the following are valid cases for multiple matches to be returned, reflecting multiple sources for data about the patient:

- Same HCID, different AAID
- Different HCID (different community, covered in previous flow)

e.

1. In step 4, the Receiving Participant returns multiple matches found. The Initiating Participant MAY choose to push the content to any or all of the matches. How it determines which is not specified.
2. The use case resumes at step 5 for each patient match the Initiating Gateway wishes to push to.

3.7.3.8 Submission references existing metadata

Informative: this flow is used for replacing documents and other purposes.

Copyright© 2020-2021 All rights reserved.

Additional precondition: the Initiating Participant has knowledge of existing metadata at the Receiving Participant. This may happen in a number of ways—see Section 3.9.10, Referencing Existing Metadata.

1. In step 5, the Initiating Participant includes the new metadata, references to the existing metadata and associations linking them, in the Document Submission.
2. The use case resumes.

3.7.3.9 Submission returns warning

1. In step 6, the Receiving Participant returns overall success for the submission but also one or more warnings.
2. The Initiating Participant takes appropriate actions—these actions are unspecified.
3. The use case ends.

3.7.4 Exception Flows

3.7.4.1 No compatible transmission option found

1. In step 1, the Initiating Participant cannot find a compatible transmission option declared by the Receiving Participant.
2. The use case ends.

3.7.4.2 Incompatible option for Metadata-Limited Document Source

1. In step 1, the Initiating Participant declares itself to be an XDR Metadata-Limited Document Source and the Receiving Participant does not declare the XDR Accepts Limited Metadata option.
2. The use case ends.

3.7.4.3 Submission returns error

1. In step 6, the Receiving Participant returns overall failure for the submission: this means at least one error and potentially warnings.
2. The Initiating Participant takes appropriate actions—these actions are unspecified.
3. The use case ends.

3.7.5 Technical Post-Conditions

The following technical post-conditions will result after the execution of this interface specification:

- Any documents or clinical items that were persisted are available for subsequent query and retrieval using the eHx Query for Documents and Retrieve Documents transactions.
- Audit logs as defined in Section 6 have been recorded.

Copyright© 2020-2021 All rights reserved.

3.7.6 QTF Interoperability

Informative: The [TEFCA QTF] will have its own directory. Assuming at this point that the eHx directory will be populated to include QTF entries, and that eHx entries will be pushed to the QTF directory, so that participants can operate as normal. QTF entries in the eHx directory will not have any of the eHx-specific options defined except for the XCDR transaction option. In addition, eHx systems declaring eHx options in the QTF directory will not have those options shown. This should not be a problem, as we have written robust handling requirements.

Further, we don't know at this point how federated systems will show up.

3.8 Use Case Flow Requirements

This table shows the required flows from the Push use case for the Initiating (I) and Receiving (R) Participants.

Rqmt #	Flow	I/R	Required to Support
CONF-016	Nominal Flow	I/R	SHALL
CONF-017	Push to implicitly federated recipient	I	SHALL, although no difference
CONF-018	Push to implicitly federated recipient	R	MAY
CONF-019	Push to explicitly federated recipient	I	SHALL
CONF-020	Push to explicitly federated recipient	R	SHALL, although may ignore if no defined recipients
CONF-021	Push to federated recipient	I/R	SHALL, although may ignore if no defined recipients
CONF-022	Patient ID already obtained	I	MAY
CONF-023	Patient ID already obtained	R	N/A
CONF-024	Patient ID is optional	I	SHALL
CONF-025	Patient ID is optional	R	MAY
CONF-026	No required patient match found	I/R	SHALL

Copyright© 2020-2021 All rights reserved.

CONF-027	Patient match returned from different community	I	SHALL
CONF-028	Patient match returned from different community	R	MAY
CONF-029	Multiple patient matches returned	I	SHALL
CONF-030	Multiple patient matches returned	R	MAY
CONF-031	Submission references existing metadata	I	SHALL: See Section 3.9.10 for allowable purposes and additional behavior requirements.
CONF-032	Submission references existing metadata	R	SHALL: See Sections 3.15 through 3.23 for processing requirements.
CONF-033	Submission returns warning	I/R	SHALL
CONF-034	No compatible transmission option found	I	SHALL
CONF-035	No compatible transmission option found	R	N/A
CONF-036	Incompatible option for Metadata-Limited Document Source	I	SHALL if a Metadata-Limited Document Source
CONF-037	Incompatible option for Metadata-Limited Document Source	R	N/A
CONF-038	No patient match found	I/R	SHALL
CONF-039	Submission returns error	I/R	SHALL

3.9 Message Requirements: Metadata Elements

The metadata passed in this transaction is constrained within this eHx specification based on anticipated use cases as follows.

Copyright© 2020-2021 All rights reserved.

CONF-040: An Initiating Participant SHALL format metadata elements as specified in [IHE ITI TF-3] 4.1 and 4.2.

CONF-041: An Initiating Participant SHALL populate metadata elements as specified in [IHE ITI TF-3] Table 4.3.1-3: Sending Actor Metadata Attribute Optionality:

- **CONF-042:** If it is utilizing the XCDR transaction option, it SHALL populate according to the "XDR DS" column, except for patientId as described below.
- **CONF-043:** If it is utilizing the XDR transaction option and does not declare itself to be a Metadata-Limited Document Source, it SHALL populate according to the "XDR DS" column.
- **CONF-044:** If it is utilizing the XDR transaction option and declares itself to be a Metadata-Limited Document Source, it SHALL populate according to the "XDR MS" column.

3.9.1 QTF Interoperability

CONF-045: An Initiating Participant, if submitting to a Receiving Participant through the QTF, SHALL populate according to the "XDR MS" column.

Informative: The [TEFCA QTF] adopts XCDR without constraints, which does not allow limited metadata.

Some of the key metadata elements are further described and constrained here:

3.9.2 XDSDocumentEntry.sourcePatientId

CONF-046: The Source Patient ID SHALL contain two parts:

- Patient Identity Assigning Authority in the form of an OID
- An identifier in the above Assigning Authority domain

Informative: The Source Patient ID represents the community identifier of the subject of care (i.e., patient) of the document from the Initiating Participant's Assigning Authority domain.

3.9.3 XDSDocumentEntry.sourcePatientInfo

CONF-047: If included, sourcePatientInfo SHOULD specify a minimum of demographics for the patient as known by the Initiating Participant, including first name, last name, date of birth and gender.

Informative: This is an optional element, required if known for the XDR Metadata-Limited Document Source.

Copyright© 2020-2021 All rights reserved.

3.9.4 XDSDocumentEntry.patientId

CONF-048: The Patient ID SHALL contain two parts:

- Patient Identity Assigning Authority in the form of an OID.
- An identifier in the above Assigning Authority domain.

Informative: The Patient ID represents the subject of care of the document (i.e., patient) from the Receiving Participant's Assigning Authority domain. This value is obtained by the Initiating Participant through some verifiable means, primarily through use of the Patient Discovery Specification.

CONF-049: An Initiating Participant utilizing the XCDR transaction option MAY omit the patientId attribute if the Receiving Participant does not declare the eHx Patient ID Required option.

Informative: The base IHE ITI standard has conflicting requirements around the patientId attribute—XCDR requires it to be included by adopting the XDR DS optionality, but [IHE ITI TF-2c] 3.80.4.1.1 allows it to be omitted. The authors of this specification are currently working with the IHE ITI Technical Committee on a Change Proposal to make patientId R2 (required if known). The above requirement is written to anticipate this clarification. This specification may need to be updated following the IHE Change Proposal process.

CONF-050: An Initiating Participant utilizing the XDR transaction option SHALL populate the patientId attribute in all of the following cases:

- The Receiving Participant declares the eHx Patient ID Required option.
- The Receiving Participant does not declare the XDR Accepts Limited Metadata option.

CONF-051: An Initiating Participant that omits the patientId attribute SHALL meet the conditions in [IHE ITI TF-2c] 3.80.4.1.1.

Informative: The above requirement requires adequate demographics to ensure identification of the patient.

CONF-052: An Initiating Participant, if populating the patientId, SHOULD populate with the subject of care of the submission set from the Receiving Participant's Assigning Authority domain.

Informative: The reason the above is a SHOULD is to allow an initiator to still Push if the patient match fails yet the patientId is required. The base IHE specification does not constrain the domain (e.g., allowing the sourcePatientId to be used in this case).

Copyright© 2020-2021 All rights reserved.

3.9.4.1 QTF Interoperability

CONF-053: An Initiating Participant, if submitting to a Receiving Participant through the QTF, SHALL populate the patientId.

Informative: The [TEFCA QTF] adopts XCDR without constraints, which means patient ID will be required, until/unless XCDR accepts our CP to make it R2.

3.9.5 XDSDocumentEntry.Hash

CONF-054: An Initiating Participant SHALL populate the Hash with the hash of the document, computed following the SHA-1 algorithm.

3.9.6 XDSDocumentEntry.Size

CONF-055: An Initiating Participant SHALL populate the Size with the actual size (in bytes) of the document.

3.9.7 XDSSubmissionSet.patientId

CONF-056: An Initiating Participant SHALL populate the Patient ID with the subject of care of the submission set from the Receiving Participant's Assigning Authority domain.

CONF-057: The Patient ID SHALL follow the same rules as defined for XDSDocumentEntry.patientId in Section 3.9.4.

3.9.8 XDSSubmissionSet.sourceId

CONF-058: An Initiating Participant SHALL populate the Source ID with its homeCommunityId.

Informative: The homeCommunityId is a globally unique identifier for a community used to assist in subsequent transactions for locating the data held by that community. homeCommunityId is structured as an OID limited to 64 characters and specified in URI syntax, for example the homeCommunityId of 2.16.840.1.113883.3.166 would be formatted as urn:oid: 2.16.840.1.113883.3.166.

3.9.9 Value Sets for Coded Attributes

CONF-059: An Initiating Participant SHALL populate coded metadata attributes according to the following [HL7 FHIR R4] value sets and binding strengths. Binding strengths are defined according to the HL7 FHIR specification: https://www.hl7.org/fhir/valueset-binding-strength.html.

CONF-060: An Initiating Participant, if populating coded metadata attributes with extended values, SHALL use value sets defined by the eHealth Exchange.

Copyright© 2020-2021 All rights reserved.

Table 1 Value sets for XDS document metadata

XDS metadata attribute	Value set	Binding strength
DocumentEntry authorRole	http://hl7.org/fhir/R4/valueset-practitioner-role.html	Preferred
DocumentEntry authorSpeciality	http://hl7.org/fhir/R4/valueset-c80-practice-codes.html	Preferred
DocumentEntry classCode	http://hl7.org/fhir/R4/valueset-document-classcodes.html	Extensible
DocumentEntry confidentiality Code	http://hl7.org/fhir/R4/valueset-security-labels.html	Extensible Informative: Note that the latest value set has many more values than historically used in the eHx. This is what FHIR DocumentReference allows.
DocumentEntry eventCodeList	http://hl7.org/fhir/R4/v3/ActCode/vs.html	Example Informative: This specification does not constrain this attribute, as it is very specific to the type of document.
DocumentEntry formatCode	http://hl7.org/fhir/R4/valueset-formatcodes.html	Extensible
DocumentEntry healthcareFacilityTypeCode	http://hl7.org/fhir/R4/valueset-c80-facilitycodes.html	Extensible
DocumentEntry practiceSettingCode	http://hl7.org/fhir/R4/valueset-c80-practice-codes.html	Extensible
DocumentEntry typeCode	http://hl7.org/fhir/R4/valueset-c80-doc-typecodes.html	Extensible

Copyright© 2020-2021 All rights reserved.

DocumentEntry mimeType	http://www.hl7.org/documentcenter/public/standards/vocabulary/vocabulary_tables/infrastructure/vocabulary/mediaType.html	Required
SubmissionSet contentTypeCode	http://hl7.org/fhir/R4/v3/ActCode/vs.html	Example

Informative: The above value sets and binding strengths are the same for the equivalent metadata elements in FHIR (see section TBD). Also note that the values for authorRole and authorSpeciality may be passed as coded values or as simple strings. The example in this specification shows both.

3.9.9.1 QTF Interoperability

Informative: At this point, the [TEFCA QTF] has not adopted any metadata requirements, so our constraints should not be a problem.

3.9.10 Referencing Existing Metadata

CONF-061: An Initiating Participant that wishes to reference existing metadata in a Document Submission MAY obtain metadata references using any of the following methods and SHALL prefer the methods in the order they are listed, unless otherwise specified:

- The Initiating Participant obtains the existing metadata references in a Query for Documents transaction.
- The Initiating Participant had specified the entryUUID for the existing metadata objects in a prior Document Submission.
- The Initiating Participant obtains the existing metadata references in an unspecified way.

Informative: Multiple scenarios, such as replacing a document, require the Initiating Participant to reference existing metadata at the Receiving Participant. An existing metadata object such as a Document Entry or a Submission Set is referenced by the entryUUID field, a key that is intended to be unique at the Receiving Participant. Typically, this value is generated internally when the object is added, but it may be explicitly supplied by the original submitter.

Copyright© 2020-2021 All rights reserved.

3.10 Message Requirements: Updating Documents

CONF-062: An Initiating Participant MAY attempt to update documents using two methods: document replacement and document appending.

Informative: This specification does not specify the conditions that trigger an Initiating Participant to submit an updated document. Further, there should not be an expectation that the Receiving Participant will always accept the update—it may apply checks and processing before accepting, especially if the source system or author differs.

3.10.1 Submitting updates to a previously submitted document

Informative: There may sometimes be a need to correct or add to a document that was previously submitted (i.e., the submitter of the original and replacement is the same author/owner). The below workflow handles the following:

- If the document was saved as is, the Receiving Participant may respond with a version clash. This behavior is specified in section 3.18.
- If the document was saved in some other format, the Receiving Participant needs to ensure replacement semantics. By explicitly submitting the new content as a replacement to a specific document identified by UUID (even if the item wasn't persisted as is), the Receiving Participant has enough information to make a proper decision.

CONF-063: An Initiating Participant that has the capability of submitting corrections to documents it has submitted SHALL specify the entryUUID for documents in all submissions and persist it for later potential corrections.

CONF-064: When updating a previously submitted document, an Initiating Participant SHOULD perform the following workflow:

- For the first attempt, the Initiating Participant specifies as the association target the entryUUID it originally submitted.
- If the first attempt fails due to a version clash (XDSRegistryDeprecatedDocumentError is returned), the Initiating Participant either abandons the update or queries to determine the latest applicable document and submits updates to that instead. Note that the latest document may have relationships to appendices or transformations.

3.10.2 Submitting updates to a discovered document

CONF-065: If an attempt to update a discovered document fails due to a version clash (XDSRegistryDeprecatedDocumentError is returned), the Initiating Participant SHOULD either

Copyright© 2020-2021 All rights reserved.

abandon the update or query to determine the latest applicable document and submit updates to that instead. Note that the latest document may have relationships to appendices or transformations.

Informative: The term “discovered document” in this context means a document that is discovered through querying the Receiving Participant. Although we do not limit updates to the original author, when performing cross-author updates there are special considerations. See the Message Requirements: Provenance section 3.14 for details.

3.10.3 Reflecting the Update in the Document

CONF-066: When replacing an HL7 CDA document, an Initiating Participant that has control over the generation of the new document SHOULD populate the relatedDocument element with a typeCode of “RPLC” and identify the prior document in the header of the new document.

CONF-067: When appending to an HL7 CDA document, an Initiating Participant that has control over the generation of the new document SHOULD populate the relatedDocument element with a typeCode of “APND” and identify the prior document in the header of the new document.

3.10.4 Reflecting the Update in the Document Metadata

CONF-068: When updating a document, an Initiating Participant SHALL reference existing metadata according to Section 3.9.10, Referencing Existing Metadata.

CONF-069: An Initiating Participant SHALL support XDS document relationships of type RPLC and APND.

CONF-070: When replacing a document, if a reference to the existing document entry can be obtained, an Initiating Participant SHALL include in the submission a RPLC association as defined in [IHE ITI TF-3] 4.2.2.2.3.

CONF-071: When appending to a document, if a reference to the existing document entry can be obtained, an Initiating Participant SHALL include in the submission an APND association as defined in [IHE ITI TF-3] 4.2.2.2.1.

3.11 Message Requirements: Other Document Relationships

Informative: An Initiating Participant MAY submit document relationships of type XFRM, XFRM_RPLC, or signs, but any expected behavior is undefined unless specified by a higher-level profile or participant agreement. See Processing Requirements, section 3.20.

Copyright© 2020-2021 All rights reserved.

Informative: An Initiating Participant MAY submit associations linking document entries to existing submission sets, but any expected behavior is undefined unless specified by a higher-level profile or participant agreement. See Processing Requirements, section 3.22.

See Section 3.9.10, Referencing Existing Metadata, for how to obtain and express references.

Informative: These relationships are not included because there has been no use case presented for them. This could change in the future.

CONF-072: An Initiating Participant SHALL NOT submit document relationships of type IsSnapshotOf.

Informative: The above requirement has been added because the IsSnapshotOf relationship is only used between an On-Demand Document Entry and a Stable Document Entry for a corresponding document that was generated. There is no reason for one participant to relate these objects for another participant.

3.12 Message Requirements: Folders

Informative: Folder semantics are defined in [IHE ITI TF-3] 4.2.1.3.

Informative: An Initiating Participant MAY submit Folders, but any expected behavior is undefined unless specified by a higher-level profile or participant agreement. See Processing Requirements, section 3.22.

See Section 3.9.10, Referencing Existing Metadata, for how to obtain and express references.

Informative: Folders are not included because there has been no use case presented for them. This could change in the future.

3.13 Message Requirements: Routing

Informative: An Initiating Participant wishing to route to additional recipients should ensure that the Receiving Participant recognizes all of them, for example, by determining the relationship to the Receiving Participant from the eHx directory. Any unknown or unreachable participant may cause the entire transaction to fail.

CONF-073: An Initiating Participant wishing to submit to a federated community using the XCDR transaction option SHALL populate the Home Community ID of the ultimate recipient in the appropriate fields as specified in [IHE ITI TF-2b] 3.41.4.1.2.2.

Copyright© 2020-2021 All rights reserved.

CONF-074: An Initiating Participant wishing to submit to a federated community using the XDR transaction option SHALL populate exactly one Submission Set intendedRecipient with the XON format including the Home Community ID of the ultimate recipient, as specified in [IHE ITI TF-3] 4.2.3.3.7.

CONF-075: An Initiating Participant wishing to submit to a non-community organization or individual SHALL populate Submission Set intendedRecipient with information identifying the ultimate recipient as specified in [IHE ITI TF-3] 4.2.3.3.7.

3.14 Provenance

Informative: In the context of cross-community sharing of clinical information, there is always a need to convey provenance to the receiver. This section specifies general provenance requirements for all clinical documents that may be shared, irrespective of whether they are shared by push, pull or subscription mechanisms. In order to ensure that the information needed is present when shared, it includes behavior requirements based on triggers prior to the sharing action.

CONF-076: The Participant SHALL follow the provenance requirements in this section for all CDA documents that may be shared according to this specification, unless overridden by another eHx Provenance specification.

Informative: The purpose of this requirement is to allow future content-specific profiles to further relax or constrain provenance as defined in this document to meet use cases.

CONF-077: The Initiating Participant SHALL follow the provenance requirements in this section when submitting clinical content to Receiving Participants.

CONF-078: The Receiving Participant SHALL follow the provenance requirements in this section when receiving clinical content from an Initiating Participant.

CONF-079: The Participant SHALL follow the provenance requirements in this section when sharing clinical content internally.

Informative: The intent of the eHx is to preserve provenance from the original source via any intermediaries to the final receiver. The above three requirements ensure that provenance is retained no matter where or when the sharing behavior occurs. For example, if a Participant internally generates a Patient Summary CCD document by aggregating content from multiple encounters, authorship needs to be retained just as in [HL7 Basic Provenance] Use Case 3 (HIE Transformation), even if the document is not immediately shared with another Participant.

Copyright© 2020-2021 All rights reserved.

CONF-080: Organizations that assemble content from Encounter Summary documents into Patient Summary Documents that may be exchanged with external organizations, SHALL retain provenance in accordance with [HL7 Basic Provenance] Use Case 3 (HIE Transformation).

CONF-081: The Participant SHALL follow all requirements in [HL7 Basic Provenance], constrained and clarified as follows (the requirement numbers cited below are within the referenced specification):

- 3.1 Basic Provenance Practices
  - CONF:1000: Constrained: SHALL apply only when the data is authored by a care providing organization or the patient. Informative: This clarifies that document content can be authored by other types of entities (e.g., payers, labs, etc.).
  - CONF:1004: Informative: In addition to CONF:1004, there are additional requirements for reconciliation in section 3.14.4 of this specification.
- 3.2.1 C-CDA Provenance Practices
  - CONF:1006 and CONF:1007: Constrained: The Provenance – Author Participation template constraints SHALL be met, and the template entry SHOULD be present. Informative: this allows for systems that can populate all the required data but not the new template ID.
- 3.3 FHIR: Omitted
- C. Provenance - Assembler Participation: Informative: Note that there is no requirement to use this template in [HL7 Basic Provenance]. We add it below.

CONF-082: The Participant SHALL use either the Provenance – Assembler Participation template from [HL7 Basic Provenance] or the Assembler Document Participant template from [HL7 Data Provenance] when generating a document entirely from existing content.

Informative: [HL7 Basic Provenance] states the templates are consistent, but there are slight differences. This allows for systems that implemented provenance to the older specification.

3.14.1 Updating a document: Appending or replacing from the same source

Informative: When updating a document, the association between the original and new document is considered an important part of provenance and is fully specified in other sections of this specification.

3.14.2 Updating a document: Replacing from a different source

Informative: This is a very constrained case. See Considerations, section 5.

CONF-083: If a Participant is generating a replacement to a document originally from another source, defined as a different community, organization or author, it SHALL:

Copyright© 2020-2021 All rights reserved.

- Use the appropriate document-level template from [HL7 Data Provenance] depending on the kind of author (e.g., Provider Generated Document With Provenance).
- Add itself to the document-level author and use the original author for those sections or entries that are not being changed.

Informative: The following two sections are for any change to existing clinical data, whether made by replacing a document the data originally appeared in or by introducing a new document that explicitly modifies existing data. They may be combined, e.g., when changing individual entries in one section and reconciling another section entirely.

3.14.3 Changing Individual Entries

CONF-084: If a Participant is generating a document that reflects changes to one or more individual entries from prior source document(s) and is not reconciling the entire section, it SHALL:

- Use the appropriate document-level template from [HL7 Data Provenance] depending on the kind of author, e.g., Provider Generated Document With Provenance.
- Use the appropriate entry-level template from [HL7 Data Provenance] for any modified entry, e.g., Observation Generated by Provider.

3.14.4 Reconciling sections

CONF-085: If a Participant wishes to perform reconciliation of one or more entire sections from prior source document(s), it SHALL follow the requirements in [IHE RECON] for Content Creator with Reconciliation Content Option, constrained and clarified as follows:

- It MAY be grouped with a Reconciliation Agent. Informative: this is a relaxation of the SHALL grouping requirement. It allows Participants to manually reconcile sections without the aid of an automated agent.
- When generating a document containing reconciled content, it SHALL conform to the content requirements in [IHE RECON] 6.3.1.D Reconciliation Content, including a Reconciliation Act in each section that has been reconciled.

Informative: For example, a patient's PCP submits a replacement document for one authored by another physician because it has incorrect information about the patient. We cover two such cases: selective correction and section-level reconciliation. Note that cross-author updates may be treated differently from ordinary submissions by the Receiving Participant, for example, additional security checks may be performed. See the Security Considerations section 5 for details.

Copyright© 2020-2021 All rights reserved.

3.15 Processing Requirements: Overall Processing

CONF-086: The Receiving Participant SHALL process the entire submission, including any federated submissions in the case of the XCDR transaction option, before returning, with one exception: if a document needs to be queued for manual matching to a patient. In this case, the Receiving Participant SHALL return a DocumentQueued warning code for each document so queued.

Informative: The base IHE transactions require full processing of the submission before returning. However, there is an XDR warning code, DocumentQueued, that appears to permit an exception to this expectation. We clarify that here.

CONF-087: If multiple exceptional conditions exist, the Receiving Participant SHALL detect and include each one in the response.

3.15.1 QTF Interoperability

Informative: In general, our behavior requirements simply add more specificity for processing and error handling. As every response or error code we return is defined in Vol 3, we do not anticipate any problems.

3.16 Processing Requirements: Patient Matching

CONF-088: If the patientId attribute is not included in the submission request, the Receiving Participant SHOULD attempt to determine the local patient to apply the submission to by matching the demographics in the sourcePatientInfo attribute.

CONF-089: If the patientId attribute is included in the submission request but is unrecognized, the Receiving Participant SHOULD attempt to determine the local patient to apply the submission to by matching the demographics in the sourcePatientInfo attribute.

CONF-090: If no local patient for the submission can be identified, the Receiving Participant MAY return an XDSUnknownPatientId error code.

Informative: We are allowing receivers to be forgiving here as they may not need the idea of a patient ID at all.

3.17 Processing Requirements: Routing

Informative: Receiving Participants that support routing to any federated communities, sub-organizations or persons will work with eHx staff or update the eHx Directory directly to ensure

Copyright© 2020-2021 All rights reserved.

each potential recipient is represented appropriately in the directory, or is included in an external directory available to participants.

CONF-091: A Receiving Participant that supports the XCDR transaction option SHALL return an XDSUnknownCommunity error if a supplied destination Home Community ID is unknown.

CONF-092: A Receiving Participant that supports the XCDR transaction option SHALL return an XDSUnavailableCommunity error if the destination community cannot be reached.

Informative: The above two requirements are not currently in XCDR. We have opened a CP with IHE to address this.

CONF-093: A Receiving Participant SHOULD return an XDSRepositoryError if the destination Home Community ID is populated in an ITI-41 submission.

Informative: A Receiving Participant that supports the XCDR transaction option receives the Home Community ID of the ultimate recipient in the appropriate fields of the ITI-80 transaction as specified in [IHE ITI TF-2b] 3.41.4.1.2.2, and routes to it (see Use Case Section 3.7.3.1, Push to federated system). The Document Submission specification does not use the variant of ITI-41 that includes HCID.

CONF-094: If the SubmissionSet.intendedRecipient is populated, the Receiving Recipient SHALL make reasonable efforts to determine whether each recipient can be notified, but MAY return success before confirming full receipt and processing by the intendedRecipients. A Receiving Recipient MAY delegate notification of some or all intendedRecipients, for example, if a federated community that can route to these recipients is specified in the XCDR fields or in intendedRecipient. If notification of an intendedRecipient is not possible, the Receiving Recipient MAY do any of the following (the Error/Warning codes are defined in [IHE ITI TF-3] 4.2.4.1, as modified by [ITI CP TBD]).

- Fail the transaction and return the code UnknownRecipient or UnavailableRecipient as an error
- Succeed and return the code UnknownRecipient or UnavailableRecipient as a warning
- Succeed silently

If the recipient is a community, the error codes XDSUnknownCommunity or XDSUnavailableCommunity should be used instead.

Copyright© 2020-2021 All rights reserved.

3.18 Processing Requirements: Persisting Content

CONF-095: If the Receiving Participant determines that document(s) and/or metadata will be persisted, it SHALL perform equivalent behavior to the XDS Document Repository as specified in [IHE ITI TF-2b] 3.41.4.1.3 and 3.41.4.2.1.1, with the following exceptions:

- Instead of or in addition to using the Register Document Set-b [ITI-42] transaction to communicate to an XDS Document Registry, it MAY perform equivalent behavior using unspecified mechanisms.
- Instead of or in addition to making persisted documents available for retrieval via the Retrieve Document Set [ITI-43] transaction, it SHALL make them available for retrieval via the Cross Gateway Retrieve [ITI-39] transaction.
- Instead of or in addition to making persisted document entries available for query via the Registry Stored Query [ITI-18] transaction, it SHALL make them available for query via the Cross Gateway Query [ITI-38] transaction.
- It SHOULD return warnings instead of errors for any nonconformant metadata fields or mime types.
- Section 3.42.4.1.3.3.2 patientId Attributes: the requirement to verify against a Patient Identity Feed [ITI-8] MAY be ignored; appropriate behavior is already covered in section 3.16.
- Other than Append and Replace associations, it MAY choose not to persist Submission Sets, other Associations and/or Folders, and if so, MAY ignore any related requirements.

Informative: The above requirements include enforcing document replacement semantics, deprecating prior versions. Also, note that the sender may explicitly specify entryUUID for objects and expect the receiver will persist them: see [IHE ITI TF-2b] 3.42.4.1.3.7.

3.19 Processing Requirements: Existing Metadata

CONF-096: The requirements for the Receiving Participant to be able to process the submission without any context (XDR: [IHE ITI TF-2b] 3.41.4.1.3.1, XCDR: [IHE ITI TF-2c] 3.80.4.1.3) do not apply when the submission contains references to existing metadata. In these cases, the Receiving Participant SHALL verify and process these references as detailed in this specification.

Informative: Both ITI-41 and ITI-80 state in Expected Actions that the receiver must be able to process the submission without any context. This is misleading, because in some cases context is required. We are submitting an IHE CP to clarify.

Copyright© 2020-2021 All rights reserved.

3.20 Processing Requirements: Document Relationships

CONF-097: A Receiving Participant SHALL accept XDS document relationships of type APND, RPLC, XFRM, XFRM_RPLC and “signs”.

Informative: This is an additional requirement over the base specs, but is needed because any system that accepts documents must be able to accept corrections to those documents, thus they must be able to accept RPLC. Further, once that relationship is supported, the others do not require any special processing, so they can be supported as well.

CONF-098: If a Receiving Participant receives a document relationship of type IsSnapshotOf, it SHALL return an XDSRepositoryMetadataError error.

Informative: If a Receiving Participant receives a document relationship of type RPLC or XFRM_RPLC and has persisted the prior document, replacement semantics are covered in section 3.18.

CONF-099: If a Receiving Participant receives a document relationship of type RPLC or XFRM_RPLC and has persisted information about the prior document in some way other than as a document, it SHALL, if possible, ensure replacement semantics are followed as appropriate, for example: remove or mark the information as superseded.

One example of the above requirement would be Public Health Reporting, where the only thing persisted was a document ID and a record of a condition. In this case, the Receiver would be required to update that record accordingly.

Informative: For the case of cross-author updates, see the Security Considerations section 5 for additional considerations.

3.21 Processing Requirements: Conflict Detection and Resolution

Informative: Other sections in this specification cover deterministic cases of conflict detection, for example, reusing the same uniqueId for a different document. This section covers the other case: when content within submitted documents may contain duplicate or conflicting information. The receiver needs to be able to detect this and return errors.

Example:

- System A submits an encounter document for a patient, and the document is persisted.
- System B submits a different encounter document representing the same real-world encounter for the patient.

Copyright© 2020-2021 All rights reserved.

- The receiving system needs to be able to detect that the encounters are the same and handle appropriately. It may reject one document, it may store both documents and/or it may correlate any information it has stored from the documents.

Example:

- System A submits to a public health system an encounter document where a patient was identified as positive for COVID-19. The document itself is not persisted; rather some minimum amount of information to allow tracking is stored.
- System B submits to the same public health system a different encounter document where the same patient was identified as positive for COVID-19.
- The receiving system needs to be able to detect that the patients are the same and handle appropriately. It may retain the separate encounters while not double-counting the positive result.

CONF-100: A Receiving Participant that persists clinical information from submitted documents in some other form SHALL be able to detect duplicate or conflicting information and return an XDSRepositoryError error if the conflict cannot be resolved successfully.

Informative: Initiating Participants that can submit corrections will always specify the entryUUID for documents they submit. See section 3.10.1.

CONF-101: A Receiving Participant that persists clinical information from submitted documents in some other form SHOULD persist the entryUUID for submitted documents, when it is specified in the submission, in order to correlate any future corrections to that document.

3.22 Processing Requirements: Folders and Submission Sets

Informative: In addition to document entries and the relationships between them, the full metadata model includes submission sets and folders, which are linked to each other and to document entries by HasMember associations. Receivers are not required to persist them.

CONF-102: A Receiving Participant MAY accept Folders, but any expected behavior is undefined unless specified by a higher-level profile or participant agreement.

CONF-103: If a Receiving Participant receives but does not support Folders, it SHALL ignore that content, process the rest of the submission and return a PartialFolderContentNotProcessed warning with the response.

Informative: The conditions for returning the “Partial…” warnings are not clear, and we are submitting a CP to IHE to clarify. The above requirements represent our interpretation, which is

Copyright© 2020-2021 All rights reserved.

that these warnings mean the entire concept is not supported, not that there was a problem with one specific submission.

CONF-104: If a Receiving Participant does not persist submission sets or their associations, it SHALL process the rest of the submission and SHOULD NOT return any error or warning with the response.

3.23 Processing Requirements: Additional Exception Checking

CONF-105: In addition to the warnings and errors specified in [IHE ITI TF-2c] 3.80.4.1.3, an XCDR Responding Gateway SHALL detect and return the additional warnings and errors specified in [IHE ITI TF-2b] 3.41.4.1.3.1 for the XDR Document Recipient.

Informative: There are differences in error checking between XDR and XCDR, which do not have an obvious purpose. We are submitting a CP to IHE to clarify this. The above requirements represent our interpretation, which is that error checking should be as similar as possible.

CONF-106: The Receiving Participant SHALL return an XDSPatientIdDoesNotMatch error if objects in a submission set do not have the same patientId as the submission set.

CONF-107: The Receiving Participant SHALL return an XDSPatientIdDoesNotMatch error if both DocumentEntry objects referenced by an Association do not have the same patientId.

CONF-108: The Receiving Participant SHOULD return an XDSRegistryMetadataError or XDSRepositoryMetadataError code with a severity of Warning if a coded value is submitted that is not within the defined value sets in Table 1 Value sets for XDS document metadata. Informative: see http://sequoiatechwg.editme.com/Topic-1537135046658 and https://healthcaresecprivacy.blogspot.com/2011/11/xdsxca-testing-of-vocabulary.html.

CONF-109: The Receiving Participant SHALL detect and return metadata errors as specified in the following sections in [IHE ITI TF-2b], as constrained below:

- 3.42.4.1.3.3.6 DocumentEntry.serviceStartTime and DocumentEntry.serviceStopTime
  - Use XDSRegistryMetadataError or XDSRepositoryMetadataError with a severity of error
- 3.42.4.1.3.5 Document Relationships
  - Excluding the check for patient identity merges
  - Replacing the term “Document Registry” with the equivalent storage mechanism for document entries.

Copyright© 2020-2021 All rights reserved.

- Excluding the behavior associated with replacement (“When the Association type is "RPLC" or "XFRM_RPLC"....”) and On-Demand snapshots (“When the Association type is "IsSnapshotOf"....”). Informative: This is persistence related, and we cover it elsewhere.

CONF-110: The Receiving Participant SHALL return an XDSRegistryDuplicateUniqueIdInMessage or XDSRepositoryDuplicateUniqueIdInMessage error if a uniqueId value was found to be used more than once within the submission.

CONF-111: The Receiving Participant SHALL return an XDSRegistryMetadataError or XDSRepositoryMetadataError error if a required metadata field is missing.

CONF-112: The Receiving Participant SHOULD return an XDSRegistryMetadataError or XDSRepositoryMetadataError error if there is any other violation of formatting rules as specified in [IHE ITI TF-3] 4.1 and 4.2. Informative: This SHOULD allows for Postel’s Law (“receive liberally”).

CONF-113: The Receiving Participant SHALL return an XDSMissingDocument error if a DocumentEntry exists in the metadata with no corresponding document.

CONF-114: The Receiving Participant SHALL return an XDSMissingDocumentMetadata error if a document is included without a corresponding DocumentEntry in the metadata.

Informative: The above requirement was prompted because the error code's description only mentions the MIME part, but this case could also happen with an unoptimized (by MTOM) binary block. We have opened a CP with IHE ITI about this.

CONF-115: The Receiving Participant SHALL return an XDSMissingDocumentMetadata error if a MIME part is attached without a corresponding Content-ID header in the metadata.

CONF-116: The Receiving Participant MAY return an InvalidDocumentContent error if the document content does not match the DocumentEntry.

CONF-117: The Receiving Participant MAY return the following error codes as defined in [IHE ITI TF-3] Table 4.2.4.1-2: XDSRegistryBusy, XDSRepositoryBusy, XDSRegistryError, XDSRepositoryError, XDSRegistryOutOfResources, XDSRepositoryOutOfResources.

3.23.1 QTF Interoperability

Informative: At this time, we don’t know if the [TEFCA QTF] will adopt any metadata constraints. If a QTF participant submits values that are outside our constraints, because we are returning a warning, not an error, this should not be a problem.

Copyright© 2020-2021 All rights reserved.

4 HL7 FHIR PUSH API DESCRIPTION

4.1 Definitions

A "resource" refers to an additional format of clinical data as it is transferred between Participants, and not as it is stored within a Participant system or specific electronic health record (EHR) system.

The following terms are defined based on which transaction option of the interface is implemented:

- The "FHIR Push Transaction" is a message exchange pattern whereby an Initiating Participant composes a message consisting of clinical, administrative and/or metadata information and sends that message to a Receiving Participant.
    - **CONF-118**: For the FHIR MHD Option, the FHIR Push Transaction SHALL correspond to the IHE ITI-65 Provide Document Bundle transaction.
    - **CONF-119**: For the FHIR Resource Option, the FHIR Push Transaction SHALL correspond to the HTTP REST create (i.e. POST) and update (i.e. PUT) methods on individual resources or in a transaction/batch Bundle.
- A "FHIR Initiating Participant" initiates a FHIR Push Transaction.
    - **CONF-120**: For the FHIR MHD Option, the FHIR Initiating Participant SHALL correspond to the Document Source actor.
    - **CONF-121**: For the FHIR Resource Option, the FHIR Initiating Participant SHALL correspond to the HTTP REST client.
- A "FHIR Receiving Participant" receives a FHIR Push Transaction.
    - **CONF-122**: For the FHIR MHD Option, the FHIR Receiving Participant SHALL correspond to the Document Recipient actor.
    - **CONF-123**: For the FHIR Resource Option, the FHIR Receiving Participant SHALL correspond to the HTTP REST server.
- A "FHIR Participant" refers to both FHIR Initiating Participants and FHIR Receiving Participants.
- **CONF-124**: When "Patient Discovery" is referenced in the context of any of the FHIR transaction options, the exact mechanism is currently out of scope but is defined as the capability of matching a patient and obtaining either a patient ID or a URL to a Patient resource that is valid at the FHIR Responding Participant. We anticipate the eHx Patient Discovery specification to be enhanced to define this for FHIR.
- **CONF-125**: When "Access Consent Policies" is referenced in the context of any of the FHIR transaction options, the exact mechanism is currently out of scope but is defined

Copyright© 2020-2021 All rights reserved.

as the capability of including a reference to a Consent resource in a FHIR request and obtaining that resource for the purposes of determining access. We anticipate the eHx Access Consent Policies specification to be enhanced to define this for FHIR.

- **CONF-126**: When "Query for Documents" is referenced in the context of any of the FHIR transaction options, the exact mechanism is currently out of scope but is defined as the capability of querying for DocumentReference or List (or any) resources for a Patient. We anticipate the eHx Query for Documents specification to be enhanced to define this for FHIR and to be renamed, e.g., to Query for Data.
- **CONF-127**: When "Retrieve Documents" is referenced in the context of any of the FHIR transaction options, the exact mechanism is currently out of scope but is defined as the capability of retrieving a document. We anticipate the eHx Retrieve Documents specification to be enhanced to define this for FHIR.

4.2 Assumptions

The following assumptions underlie this interface specification:

- The primary expected use for the FHIR MHD Option is that documents are formatted as XML data following the HL7® Clinical Document Architecture (CDA®) standard (used with permission), but nothing precludes this interface from being used to submit other kinds of documents, such as Adobe Portable Document Format (PDF) files or images.
- The patient to whom the document(s) or resource(s) pertain:
    - Is registered at one or more facilities at the Initiating Participant.
    - Has provided consent to share his or her clinical data, or such consent is not required by the business case under which the FHIR Push Transaction is occurring; if consent is needed, the mechanism for providing this consent is the subject of the Access Consent Policies specification document.
- This transaction is between one client and one server, and any coordination between servers, for example to manage shared identities (see https://www.hl7.org/fhir/managing.html#registries), is out of scope.

4.3 Triggers

The FHIR Initiating Participant, based on a human decision or an automated workflow, wants to submit clinical information about a patient to a Receiving Participant.

4.4 Transaction Standard

CONF-128: The FHIR Push Transaction transaction is defined with three transaction options:

Copyright© 2020-2021 All rights reserved.

- **FHIR MHD Option**: This utilizes the IHE ITI-65 Provide Document Bundle transaction for the IHE Mobile access to Health Documents (MHD) profile, defined in [IHE ITI TF-2] 3.65. It supports the submission of documents and related metadata roughly equivalent to XDR, and this specification adds capabilities equivalent to XCDR.
- **FHIR Resource Option**: This utilizes the base HTTP RESTful POST and PUT transactions to support submission of individual FHIR resources or Bundles. This option may only be used for the use cases defined in this specification or in concert with an appropriate content-based profile.
- **FHIR Messaging Option**: This utilizes FHIR Messaging to send FHIR content. This option would typically only be used in concert with an appropriate content-based profile. Note that this option is only a placeholder; it is not currently in scope.

The locations and versions of these specifications, as well as other foundational standards for this transaction, are listed in Section 1.11, “Related Documents”.

Informative: Virtually all FHIR usage in practice adopts one or more profiles that are specific to the use case. This specification is intended to complement those and not to conflict with them.

CONF-129: A FHIR Participant MAY support different combinations of transaction options in either direction. Informative: For example, the FHIR MHD Option as a FHIR Initiating Participant and all options as a FHIR Receiving Participant.

CONF-130: If a FHIR Participant supports multiple transaction options, it SHALL ensure they function identically except where the transactions inherently differ. Informative: For example, any persistence or error handling implemented for one option must be the same on the other option.

CONF-131: FHIR Participants supporting the FHIR MHD Option MAY support the following IHE profile options:

IHE Actor	Supported Options
MHD Document Source	Comprehensive Metadata UnContained Reference
MHD Document Recipient	Comprehensive Metadata XDS on FHIR* UnContained Reference

Copyright© 2020-2021 All rights reserved.

*The XDS on FHIR Option groups the MHD Document Recipient with an XDS Document Source. While FHIR Participants MAY support this option, no expected behaviors are required.

CONF-132: Unless otherwise specified, FHIR Participants SHALL follow all requirements for the respective IHE actors.

CONF-133: A FHIR Initiating Participant implementing any FHIR transactions SHALL be grouped with an IHE ATNA Secure Node or Secure Application actor.

CONF-134: A FHIR Initiating Participant implementing any FHIR transactions SHALL be grouped with an IHE CT Time Client actor.

CONF-135: A FHIR Receiving Participant implementing any FHIR transactions SHALL be grouped with an IHE ATNA Secure Node actor.

CONF-136: A FHIR Receiving Participant implementing any FHIR transactions SHALL be grouped with an IHE CT Time Client actor.

Informative: The MHDS profile (TBD link) provides higher level groupings of actors of related IHE FHIR profiles to define an HIE infrastructure. This profile is not required by this specification but is referenced informationally.

4.4.1 Must Support

CONF-137: This specification adopts the definition of Must Support from [US Core STU3], section 2.1.1.4 (http://hl7.org/fhir/us/core/general-guidance.html#must-support), with the following change:

- The first sub-bullet should read: “US Core Responders SHALL be capable of populating that data element…”

4.5 Additional Options

This interface specification defines the following additional options. See the Operational Considerations section of this document for Directory considerations.

eHx Actor/Transaction Option	Supported eHx Options
FHIR Receiving Participant / FHIR MHD Option	Patient Required

Copyright© 2020-2021 All rights reserved.

- **Patient Required Option**: A FHIR Receiving Participant that declares this option indicates that it requires a patient resource known to it to be included in the submission. This is not needed with the “Comprehensive Metadata” or “XDS on FHIR” options, as the patient is included with full metadata.

4.6 Technical Pre-conditions

The following technical pre-conditions exist for this interface specification:

- The clinical information being transmitted pertains to a specific, single patient.

4.7 FHIR Usage Compared to SOAP/CDA

This section introduces FHIR and identifies what is similar and what is different from the historical mechanisms of SOAP web services and CDA for exchanging clinical content. It provides guidance for implementers in choosing how to incrementally add support for FHIR, as well as how to best take advantage of its benefits and avoid its pitfalls.

In many ways FHIR is an evolution of these older mechanisms, so you will see many things that are not just familiar, but nearly identical.

Documents as **the window into a patient’s state**

Typically, a patient receives care by providers making use of an EHR. Within this EHR, there are ways of accessing and updating information about the patient: recording vitals, adding encounters, updating symptoms and diagnoses, prescribing medications. A CDA document allows this EHR to share **a snapshot in time** of some of this information. FHIR allows the same: using FHIR, an EHR can share the exact same CDA document or a FHIR document that contains the same information.

CDA Entries vs. FHIR Resources

Within a CDA document, **a unit of information like a given allergy** is **contained in an Entry**. The entry is part of the document; it cannot exist outside it. However, there are hints of the real-life object it represents: the same allergy can be represented in multiple documents over the patient’s life, even changing state as it goes. The EHR can use the same entry identifier to let consumers know this is the same logical object. The consumer of these documents can then see the change over time of the allergy through these successive snapshots.

Copyright© 2020-2021 All rights reserved.

However, the consumer can't simply ask "What's the state of this allergy today?", or access the allergy outside the context of a document. With FHIR, they usually can, because the **logical allergy object is represented as a FHIR Resource** (https://www.hl7.org/fhir/resource.html), an independently addressable object with its own URL at a server. A resource may change over time, and its server may retain historical versions.

This doesn't completely remove complexity, however. Consumers still have to deal with the fact that the same allergy can be represented by multiple systems, each having its own view into it. Further, even within a given system, there isn't a guarantee that the EHR or HIE will have harmonized all information about that allergy into a single resource. It could have, for example, simply taken all the documents it had about a patient, and from each entry generated a separate FHIR resource, each a single snapshot in time, unrelated to any other resource.

But apart from that edge case, **when a consumer receives a FHIR resource, they typically receive both a snapshot of information plus the ability to track the information over time**. They could even attempt to correct the information if it is wrong. This is a key differentiator for FHIR.

SOAP Web Services vs. FHIR RESTful API and FHIR Messaging

SOAP web services, as specified by eHx, bind a request-response pattern to the HTTP POST transaction, and use WS-Addressing headers to tell the endpoint how to route the message. FHIR also supports a messaging layer (https://www.hl7.org/fhir/messaging.html), but most implementers use what is called the "RESTful" API (https://www.hl7.org/fhir/http.html), which uses native HTTP methods of GET, POST, PUT, etc. between client and server in a simpler and more intuitive way. This specification primarily uses the RESTful API.

Uses supported by both **mechanisms**

You can accomplish many of the same use cases with either SOAP Web Services or RESTful FHIR, but certain uses are simpler and more efficient using FHIR:

- Discover a patient by comparing demographics.
- Exchange documents to use the information in a larger clinical workflow: patient care, clinical decision support, adverse event reporting, etc.

Track the state of that workflow.

- Exchange a self-contained clinical document.
- Exchange a document that includes references to source information:
    - In CDA:

Copyright© 2020-2021 All rights reserved.

        - Information about the source of information may be captured in author person/organization.
        - Correlations to common source information may be inferred by entry identifiers.
        - An entry can reference another entry in another document to show provenance, but this mechanism is not often used and is not well understood.
    - In FHIR, a resource can reference another resource, and this is typical usage.
- Receive a document and store it as is.
- Receive a document and extract and store selected information.
- Receive a document and use references to source information to create a web of information about a patient:
    - In CDA, retain document IDs, the systems you obtained them from, and entry IDs for future correlation.
    - In FHIR, retain the resource URL

Uses made possible by FHIR

Some uses can only be accomplished using FHIR:

- Exchange information without the context of a document, e.g. an Encounter, Observation, Condition, or Task.
- Save references to individually addressable clinical resources to create a web of dynamically updated information about a patient.
- Exchange information where some sub-elements are only included by reference, and pull that referenced information from other sources.
- Collaboratively build and maintain a structure of related resources on a server.

Uses made more feasible by FHIR

Finally, there are uses that are technically possible with either mechanism (or HL7 V2 for that matter), but become much more feasible due to the lighter footprint and technical flexibility of FHIR. For example:

- Exchange information or participate in a clinical workflow from a smart phone.
- Exchange information from the context of an app running in an EHR (e.g., SMART-on-FHIR).
- Send clinical information from a device to a server.

Roadmaps to FHIR adoption for existing **participants**

Copyright© 2020-2021 All rights reserved.

For participants who are going to add FHIR capabilities to existing SOAP/CDA capabilities, how they prioritize rolling out FHIR features is ultimately up to them, but here is some guidance:

- Beyond the FHIR specification, there's a lot of good implementer information in the FHIR Community: https://chat.fhir.org/ .
- Use tools like Postman, HAPI, Touchstone and Inferno to implement and test.
- Participate in FHIR Connectathons.
- Consult with your current and potential partners on what FHIR capabilities they will be supporting and in what priority.
- See the IHE MHDS profile (https://wiki.ihe.net/index.php/Mobile_Health_Document_Sharing_(MHDS)) for guidance on how all the related IHE FHIR actors can and should be deployed in an HIE. This can assist with vendor and product selection.
- While FHIR documents are theoretically equivalent to CDA documents in terms of what information is available, keep in mind that some systems make extensive use of rendering capabilities for CDAs and C-CDAs. For this reason, consider continuing to support the CDA format even after adding support for FHIR documents.
- When choosing whether to make a clinical note available as a DocumentReference or a DiagnosticReport, consult this guidance from US Core: http://hl7.org/fhir/us/core/clinical-notes-guidance.html#fhir-resources-to-exchange-clinical-notes .
- Understand that simply enabling fine-grained FHIR resources with the bare minimum capabilities (e.g. transforming each CDA entry to an unrelated, immutable FHIR resource rather than harmonizing) may not yield partner satisfaction. People have been sold a vision of FHIR that implies harmonized, living resources.

4.8 Federation in FHIR

Informative: The concept of federation is relatively underspecified in FHIR at this time. The notion of "community" is used by numerous IHE profiles to enable complex, large-scale heterogeneous networks. See [IHE ITI TF-1] E.9 "XCA Integration with XDS and non-XDS communities" for a number of examples of federated deployments enabled by XCA. FHIR does not have an explicit analog for a community or for Home Community ID (HCID). The authors of this specification are advancing an IHE ITI work item (https://github.com/IHE/IT-Infrastructure/issues/142) to address federation in its FHIR-based profiles, and this specification will follow its design. In addition, HL7, through the FAST workgroup, is working through some federation use cases and mechanisms (the X- headers explained below). This specification does not utilize all possible mechanisms; while it doesn't disallow them, their behavior is left unspecified.

Copyright© 2020-2021 All rights reserved.

A key enabler of federation is a comprehensive directory that enables any participant to find any other, discover its HCID and electronic endpoints, and in the case of federated participants not reachable directly, to find a "parent" participant by which they can be reached (see 4.9.2.1 Push to federated system). eHealth Exchange has such a directory, and this specification leverages its capabilities. Note that these parent-child relationships don't have to be exclusive, for example: if a directory serves multiple exchanges, a given community might be reachable via different parents.

Let us consider federated cases for the FHIR Push Transaction.

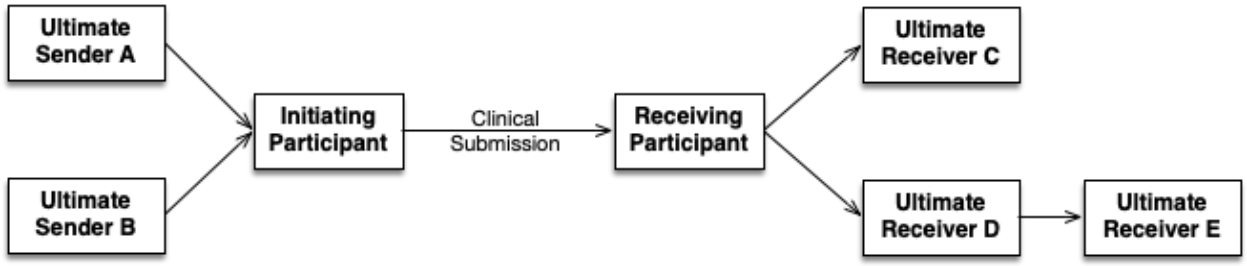

Figure 3 Federated FHIR Push Transaction

In the above diagram, the Initiating and Receiving Participants are able to send and receive submissions on their own behalf or on behalf of their child organizations. The Initiating Participant is the parent of communities A and B, the Responding Participant is the parent of communities C and D, and community D is the parent of community E. The sender and receiver roles are left generic to illustrate different ways federation and routing may be used in a Push setting.

4.8.1 Federated senders

In submissions, the identity of the ultimate sender may be passed in different places, supporting different uses. None are used by this specification.

- Transport layer: in the HTTP X-Originator header. See the ONC/FAST Hybrid/Intermediary Exchange (http://build.fhir.org/ig/HL7/fhir-exchange-routing-ig/branches/main/index.html)
- Security layer: in the OAuth client assertion. See the eHx Authorization Framework. Depending on the receiver's security model the receiver (Resource Server in OAuth parlance) or a separate Authorization Server may evaluate it as part of its access decision.
- Application layer:
    a. In an MHD Submission Set in List.extension:sourceId (see http://build.fhir.org/ig/IHE/ITI.MHD/branches/master/StructureDefinition-

Copyright© 2020-2021 All rights reserved.

IHE.MHD.Comprehensive.SubmissionSet-definitions.html#List.extension:sourceId).

b. Inferred by numerous references in a submission, for example, List.source, Patient.managingOrganization, DocumentReference.sourcePatientInfo, Resource.meta.source.

Examples:

1. Receiver uses identity of ultimate sender as part of its access decision: Example: Ultimate Receiver C limits what Ultimate Sender A can see.
2. **Initiating Participant uses identity of ultimate sender to route response**: Example: Ultimate Sender B uses asynchronous REST or messaging to submit and Initiating Participant uses the sending HCID to route the response back.
3. Receiver logs identity of ultimate sender for auditing: Any hop can do this.

4.8.2 Federated receivers

Likewise, the following mechanisms for identifying federated receivers are possible. All are supported by this specification except for X-Destination.

1. The ultimate receiver, in addition to being reachable via its federated parent, has its own FHIR server that is independently reachable as a distinct Receiving Participant. It has a unique endpoint in the directory that points directly to its FHIR server.
2. The ultimate receiver is reachable electronically, but requires mediation, e.g., to cross a security boundary, to transform from FHIR to XDR, Direct, internal messaging system, etc. Its parent uses **URL-based routing**, so the ultimate receiver has a unique endpoint in the directory that points to the parent's FHIR server, for example: HCID 1.2.3 base URL = mygateway.org/hcid/1.2.3/. See use case flow 4.9.2.1 Push to implicitly federated system.
3. The ultimate receiver is reachable electronically but requires mediation. Its parent uses **application-based routing**, so the ultimate receiver has no endpoint in the directory. This lets clients know they need to talk to a parent and pass the destination address at the application level. See use case flow 4.9.2.2 Push to explicitly federated system.
4. (not used by this specification) The ultimate receiver is reachable electronically, but requires mediation. Its parent cannot use **URL-based routing** for some reason, and the desired transaction is an individual resource create or update, not an MHD Provide Document Bundle. The sender uses the X-Destination HTTP header to route to the ultimate recipient.

Note that cases 1 and 2 do not require the sender to be aware of any federation, nor to do anything special to reach the receiver it wants.

Copyright© 2020-2021 All rights reserved.

In all of these examples, the sender may take advantage of parent-child relationships in the directory to widen or narrow their scope of their transaction. This is typically more useful in a Pull setting, for example, searching for a patient in a state HIE vs. a hospital group.

4.9 System-Level Use Case

The following use case maps the desired capabilities to the FHIR mechanisms at a system level.

4.9.1 Use Case Steps – "Nominal Flow"

Note: While the entire workflow is described here, the usage of the eHx Directory and the eHx Patient Discovery transaction are not detailed in this specification.

1. This use case begins when the Initiating Participant looks up a Receiving Participant that it wishes to push clinical information to in the eHx Directory.
2. The Initiating Participant obtains the Receiving Participant's base FHIR endpoint from the Directory and optionally its Capability Statement. The Initiating Participant may examine the capabilities declared by the Receiving Participant on eHx Patient Discovery and its desired FHIR Push Transaction.
3. The Initiating Participant sends a Patient Discovery request to the Receiving Participant to attempt to match the patient by demographics.
4. The Receiving Participant compares the demographics to its known patients and returns a Patient Discovery response to the Initiating Participant. The response contains a single patient match: a reference to a patient resource as known by the Receiving Participant.
5. The Initiating Participant sends a FHIR Push Transaction request about this patient to the Receiving Participant.
6. The Receiving Participant receives the information, processes it as appropriate and returns a FHIR Push Transaction response.

4.9.2 Alternate Flows

4.9.2.1 Push to implicitly federated recipient

f. In step 2, the Receiving Participant found in the directory has a FHIR endpoint that resolves to a federated façade or intermediary Participant, which is in the directory.
g. The use case resumes at step 2, with the following changes:
h. In steps 3 and 5 there is no difference from the Initiating Participant's perspective,
i. In steps 4 and 6, the intermediary Participant forwards the requests to and receives the responses from the actual Receiving Participant, using unspecified means.

Copyright© 2020-2021 All rights reserved.

4.9.2.2 Push to explicitly federated recipient

1. In step 2, the Receiving Participant found in the directory does not have a FHIR endpoint. However, the Initiating Participant locates a “Parent Participant” in the directory that does.
2. The use case resumes at step 2, with the following changes:
3. The Parent Participant replaces the original Receiving Participant as the system the Initiating Participant interacts directly with.
4. In step 5, the Initiating Participant includes the Home Community ID for the target (i.e., child) Receiving Participant in the FHIR Push Transaction request. See section 4.12.4 Routing.
5. In step 6, the Parent Participant routes the request to the target Receiving Participant using unspecified mechanisms.

4.9.2.3 Patient already obtained

1. In step 3, the Initiating Participant already has the patient reference.
2. The use case resumes at step 5.

4.9.2.4 Patient is optional

1. In step 3, a patient reference is optional for the desired FHIR Push Transaction. (for MHD, see Section 4.12.1.2, DocumentReference.subject).
2. The Initiating Participant chooses to execute one of the following subflows:
    a. Skip the patient match: The use case resumes at step 5.
    b. Attempt the patient match anyway. The use case resumes. If the match fails for any reason, the Initiating Participant MAY end the use case or resume at step 5.

4.9.2.5 No required patient match found

1. In step 4, the Receiving Participant returns no match found, and the patient reference is required.
2. The Initiating Participant may attempt to obtain the patient resource by manual means.
3. If the patient resource cannot be obtained, the Initiating Participant chooses to execute one of the following subflows:
    a. Resume the use case at step 5, supplying the required Patient reference as a logical reference to a business identifier, using its own Patient.identifier.
    b. End the use case.

4.9.2.6 Patient match returned from different Receiving Participant

1. In step 4, the Receiving Participant returns a patient reference from a different Receiving Participant.
2. The Initiating Participant may look up that Participant in the eHx Directory and may examine the capabilities declared by the Receiving Participant.

Copyright© 2020-2021 All rights reserved.

3. The use case resumes at step 5, with the following changes:
    a. The new Receiving Participant replaces the original Receiving Participant.

4.9.2.7 Multiple patient matches returned

1. In step 4, the Receiving Participant returns multiple patient references, which may reflect multiple identities for the same patient or multiple patients who need to be further disambiguated. The Initiating Participant MAY choose to push the content to any or all of the matches. How it determines which is not specified.
2. The use case resumes at step 5 for each patient instance the Initiating Gateway wishes to push to.

4.9.2.8 Submission updates existing resource – base FHIR

Additional precondition: The Initiating Participant has knowledge of an existing resource instance at the Receiving Participant. See Section 4.11.4 Obtaining references to resource instances.

1. In step 5, the Initiating Participant sends an “update” (i.e., PUT) FHIR Push Transaction to the resource instance.
2. The use case resumes.

4.9.2.9 Submission updates existing DocumentReference – MHD

Additional precondition: The Initiating Participant has knowledge of an existing DocumentReference instance at the Receiving Participant, either according to Section 4.11.4 Obtaining references to resource instances, or by having previously submitted the DocumentReference.

1. In step 5, the Initiating Participant includes the new DocumentReference (and accompanying document), which references the existing DocumentReference, in an MHD ITI-65 Transaction.
2. The use case resumes.

4.9.2.10 Submission updates existing Folder – MHD

Additional precondition: The Initiating Participant has knowledge of an existing MHD Folder List instances at the Receiving Participant. See Section 4.11.4 Obtaining references to resource instances.

1. In step 5, the Initiating Participant includes the updated List with a PUT method in an MHD ITI-65 Transaction.
2. The use case resumes.

Copyright© 2020-2021 All rights reserved.

4.9.2.11 Submission returns warning

1. In step 6, the Receiving Participant returns overall success for the submission but also one or more warnings.
2. The Initiating Participant takes appropriate actions—these actions are unspecified.
3. The use case ends.

4.9.3 Exception Flows

4.9.3.1 No compatible transaction option found

3. In step 1, the Initiating Participant cannot find a compatible transaction option declared by the Receiving Participant.
4. The use case ends.

4.9.3.2 Submission returns error

4. In step 6, the Receiving Participant returns overall failure for the submission: this means at least one error and potentially warnings.
5. The Initiating Participant takes appropriate actions—these actions are unspecified.
6. The use case ends.

4.9.4 Technical Post-Conditions

The following technical post-conditions will result after the execution of this interface specification:

- Any documents or clinical items that were persisted as FHIR resources are available for subsequent query and retrieval using the eHx Query for Documents and Retrieve Documents transactions.
- Audit logs as defined in Section 6 have been recorded.

4.10 Use Case Flow Requirements

This table shows the required flows from the Push use case for the FHIR Initiating (I) and Receiving (R) Participants.

Rqmt #	Flow	I/R	Required to Support
CONF-138	Nominal Flow	I/R	SHALL
CONF-139	Push to implicitly federated recipient	I	SHALL, although no apparent difference

Copyright© 2020-2021 All rights reserved.

CONF-140	Push to implicitly federated recipient	R	MAY
CONF-141	Push to explicitly federated recipient	I	SHALL if FHIR MHD option
CONF-142	Push to explicitly federated recipient	R	SHALL if FHIR MHD option, although may ignore if no defined recipients
CONF-143	Patient already obtained	I	MAY
CONF-144	Patient already obtained	R	N/A
CONF-145	Patient is optional	I	SHALL
CONF-146	Patient is optional	R	MAY
CONF-147	No required patient match found	I	SHALL
CONF-148	No required patient match found	R	MAY
CONF-149	Patient match returned from different Receiving Participant	I	SHALL
CONF-150	Patient match returned from different Receiving Participant	R	MAY
CONF-151	Multiple patient matches returned	I	SHALL
CONF-152	Multiple patient matches returned	R	MAY
CONF-153	Submission updates existing resource – base FHIR	I/R	SHALL if FHIR MHD option
CONF-154	Submission updates existing DocumentReference – MHD	I/R	SHALL if FHIR MHD option
CONF-155	Submission updates existing Folder – MHD	I/R	MAY if FHIR MHD option
CONF-156	Submission returns warning	I/R	SHALL

Copyright© 2020-2021 All rights reserved.

CONF-157	No compatible transmission option found	I	SHALL
CONF-158	No compatible transmission option found	R	N/A
CONF-159	Submission returns error	I/R	SHALL

4.11 Initiating Participant Common FHIR Requirements

CONF-160: FHIR Initiating Participants SHALL follow the requirements and guidance in [IHE ITI TF-2] Appendix Z for the "client actor", as well as any general requirements and guidance, e.g., Section Z.9 FHIR Data Types.

4.11.1 Obtaining references to HTTP URLs

Informative: As described in [FHIR R4] 3.1.0 RESTful API (https://www.hl7.org/fhir/http.html#3.1.0), clients perform create and update transactions on resource type and resource instance URLs respectively. This section describes how to obtain those URLs.

CONF-161: A FHIR Initiating Participant SHALL obtain resource type URLs as follows:

- Obtain the URL for the FHIR Receiving Participant from the eHx directory. This URL may have query parameters included, giving it the following form:

[base] {?parameter=value{&parameter=value} ...}

- Append the resource type to the base URL, while retaining any query parameters:

[base]/[type] {?parameter=value{&parameter=value} ...}

4.11.2 Creating resources

Informative: The "create" interaction is an HTTP POST command on the FHIR Receiving Participant's FHIR base URL as described in [FHIR R4] 3.1.0.8 create (see https://www.hl7.org/fhir/http.html#create). See section 4.16.1 for FHIR Receiving Participant requirements. Note that there are variations in what may be created—the exact resource POSTed may differ from what was persisted, and may not have been persisted as a resource at all. See [FHIR R4] 3.1.0.13 Transactional Integrity (https://www.hl7.org/fhir/http.html#transactional-integrity) and 7.14 Variations between Submitted data and Retrieved data (http://hl7.org/fhir/R4/updates.html#7.14).

Copyright© 2020-2021 All rights reserved.

CONF-162: A FHIR Initiating Participant SHOULD process a create response for an individual POST or a POST included in a POSTed transaction Bundle as follows:

- If the response status is 201, check for a resource in the response body. This could reflect that the resource as persisted was modified from what was POSTed or merged with an existing resource.

CONF-163: A FHIR Initiating Participant that needs the ability to correct a submitted resource SHOULD process a create response for an individual POST or a POST included in a POSTed transaction Bundle as follows:

- If the response status is 201, persist the URL of the resource for potential corrections.
- If a version is not returned in the Location or ETag, persist the Last Modified value of the resource (either from the Last-Modified header or meta.lastUpdated) for potential corrections.

Informative: If a submitted resource was stored in some other way than as an addressable resource, it will not be able to be retrieved, but it may be able to be updated. For DocumentReference in the FHIR MHD option. See section 4.12.2.4. For all other resources, see section 4.11.3.

Informative: This specification does not currently make use of the trial use Conditional Create interaction (see FHIR section 3.1.0.8.1: https://www.hl7.org/fhir/http.html#ccreate). While it handles the case where a submitted resource matches an existing one, we believe it is unnecessarily constraining on server behavior, as it does not consider merging the resources, a choice we think will be more applicable in cross-community submissions. Further, it requires additional client behavior.

4.11.3 Updating resources

CONF-164: A FHIR Initiating Participant updating a FHIR Resource via a PUT method on an individual resource or a PUT method within a transaction/batch Bundle SHALL attempt to ensure it has added its updates to the most recent version of the resource using the following mechanisms:

- Attempt to fetch the resource immediately before updating.
- If the prior version of the resource has a version identifier, include it in an If-Match header (see https://www.hl7.org/fhir/http.html#concurrency and https://tools.ietf.org/html/rfc7232#section-2.4 and https://tools.ietf.org/html/rfc7232#section-3.1).

Copyright© 2020-2021 All rights reserved.

- If the prior version of the resource has a Last-Modified value and it is being updated with an individual PUT, include the value in an If-Unmodified-Since header (see https://tools.ietf.org/html/rfc7232#section-3.4—note that this mechanism is not mentioned in [FHIR R4] and is not supported within a transaction Bundle. See https://jira.hl7.org/browse/FHIR-31927.
- If the FHIR Initiating Participant is making a correction or update to a resource it had originally created yet the resource cannot be retrieved, it SHOULD submit the update using the original resource URL returned on create. If the server had persisted information from the resource in any way, it may be able to make the correction.

Informative: The above mechanisms attempt to ensure that one client does not inadvertently overwrite the updates of another (e.g., writing stale data or incurring a race condition). However, some servers that do not support these mechanisms may handle this instead by examining resource content, not requiring client coordination. See section 4.16.2 Updating resources.

4.11.3.1 Update MHD-related resources only using MHD

CONF-165: A FHIR Initiating Participant SHALL NOT update the following document-related resources by means of PUT method on an individual resource or a PUT method within a transaction/batch Bundle:

- A DocumentReference
- A Binary or a Bundle referenced by a DocumentReference
- A List profiled as an MHD Submission Set

CONF-166: A FHIR Initiating Participant SHALL NOT update a List profiled as an MHD Folder by means of a PUT method on an individual resource or a PUT method within a transaction/batch Bundle that is not profiled as an MHD ITI-65 transaction.

Informative: Documents and related metadata may only be updated using the FHIR MHD option. See section 4.12. Note that updating documents through MHD does have deterministic update collision avoidance.

4.11.4 Obtaining references to resource instances

CONF-167: When obtaining resource instances at a FHIR Receiving Participant to update or to include by reference in a FHIR Push Transaction**,** A FHIR Initiating Participant MAY obtain these resources using any of the following methods and SHALL prefer the methods in the order they are listed, unless otherwise specified:

- For a Patient resource:

Copyright© 2020-2021 All rights reserved.

    - The FHIR Initiating Participant obtains the resource as a RESTful server-based URL in a Patient Discovery transaction or equivalent.
    - The FHIR Initiating Participant obtains the resource reference as a business identifier known to the FHIR Receiving Participant in a Patient Discovery transaction or equivalent.
- For any other resource:
    - The FHIR Initiating Participant obtains the resource as a RESTful server-based URL in a Query for Documents transaction or equivalent.
    - The FHIR Initiating Participant obtains the resource reference as a business identifier known to the FHIR Receiving Participant in a Query for Documents transaction or equivalent.
- For any resource, the FHIR Initiating Participant obtains the resource in an unspecified way.

Informative: Multiple Push scenarios, such as replacing a document, require the FHIR Initiating Participant to reference existing resources at the FHIR Receiving Participant.

4.11.5 Populating resource references

Informative: In a FHIR Push Transaction, submitted FHIR resource instances typically contain references to other resources (https://www.hl7.org/fhir/references.html). In a common case, Observation/123 refers to Patient/456, and both resources are resolveable at the same RESTful server by adding the same base URL. When submitting resources from one community to another, any included references must be understood by the receiver. The following requirements ensure that submitted references are one of the following:

- A resolvable resource at the recipient
- A business identifier known by the recipient
- A resolvable resource at a server the recipient can reach
- A resource included in the same context, e.g., a resource in the same Bundle or a contained resource

Note that the constraints on business identifiers only apply to their use as resource references. That is because a reference is the only way to resolve a resource. By contrast, it would be ok to submit a Patient containing multiple business identifiers that are not known to the recipient.

CONF-168: When submitting a resource reference (https://www.hl7.org/fhir/references.html) that is a business identifier, a FHIR Initiating Participant SHALL ensure that the identifier is one that could be known or discoverable by the FHIR Receiving Participant, for example:

- A known third-party identifier, e.g., NPI for a practitioner

Copyright© 2020-2021 All rights reserved.

- A business identifier local to the FHIR Receiving Participant

Informative: FHIR Initiating Participants should be aware that referencing resources by business identifier can lead to indeterminate results at the FHIR Receiving Participant, e.g., a query response with 0, 1 or more results.

CONF-169: When submitting a resource reference (https://www.hl7.org/fhir/references.html) that is a URL that does not resolve locally (i.e., in the same Bundle or in a Contained resource), a FHIR Initiating Participant SHALL ensure that the reference is one that could be retrieved by the FHIR Receiving Participant, for example:

- A resolvable resource hosted at the FHIR Receiving Participant
- A resolvable resource hosted at a server listed in the eHealth Exchange directory
- A canonical URL to a publicly available profile

Informative: The requirements for contained resources (https://www.hl7.org/fhir/references.html#contained) make clear that these are not simply a way to store arbitrary content. Rather, they can only be used for including resources that are part of the containing resource's definition, and the contained resource must be referenced from the containing resource. For example: Organization.contained.Endpoint referenced by Organization.endpoint.

4.11.6 Populating resource references in a transaction Bundle

Informative: Within a submitted transaction Bundle, each entry represents a resource that is to be acted upon by an HTTP verb. This specification covers only created (i.e., POST) or updated (PUT) entries.

- The rules for populating Bundle.entry.fullUrl for POSTed entries are not entirely clear in the definition for Bundle; they are made clearer in [FHIR R4] section 3.1.0.11.2 Transaction Processing Rules (https://www.hl7.org/fhir/http.html#trules). Specifically, fullUrl may be omitted; the server only needs it as a temporary identifier, in order to resolve in-bundle references from other resources, and it is ignored otherwise. If the client passes a UUID or OID value or passes a resolvable URL to the resource at its source, it should not assume the server will persist this value in any way.

Informative: There are two ways to reference a resource by business identifier within a transaction Bundle. By logical reference:

```
<patient>
  <identifier>
    <system value="http://hl7.org/fhir/sid/us-ssn" />
    <value value="000111111" />
```


Copyright© 2020-2021 All rights reserved.

```
  </identifier>
</patient>
```

Or by conditional reference (https://www.hl7.org/fhir/http.html#trules):

```
<patient>
  <reference value="Patient?identifier=http://hl7.org/fhir/sid/us-ssn|000111111"/>
</patient>
```

The processing rules are different in base FHIR—see http://hl7.org/fhir/bundle.html#references as well as the proposed update to this text: https://jira.hl7.org/browse/FHIR-29271?focusedCommentId=183020&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-183020. A logical reference carries no expectation that the server will attempt to resolve it, while a server will try to resolve a conditional reference to exactly one result, failing the transaction if it cannot.

This specification adds the requirement for servers to resolve logical business identifiers, so there is effectively no distinction. Still, if strict atomic validation and processing is desired, conditional references should be used.

4.11.7 Value Sets for Coded Attributes

Informative: FHIR, through its base types and profiles, allows for explicitly defining value sets and binding strengths for all coded attributes. This section builds on and constrains those underlying requirements. When underlying value sets are already Required and hence not further constrainable (e.g., DocumentReference.docStatus), they are not listed here.

CONF-170: A FHIR Initiating Participant SHALL populate coded metadata attributes according to the following [HL7 FHIR R4] value sets and binding strengths. Binding strengths are defined according to the HL7 FHIR specification: https://www.hl7.org/fhir/valueset-binding-strength.html .

CONF-171: A FHIR Initiating Participant, if populating coded metadata attributes with extended values, SHALL use value sets defined by the eHealth Exchange.

Table 2 Value sets for FHIR resources

FHIR Resource Elements	Value set and binding strength
DocumentReference.type	Extensible: http://hl7.org/fhir/R4/valueset-c80-doc-typecodes.html US CORE alternate: http://build.fhir.org/ig/HL7/US-Core-R4/ValueSet-us-core-documentreference-type.html .

Copyright© 2020-2021 All rights reserved.

	Required binding, links to same LOINC document types and adds nullFlavor UNK.
DocumentReference.category	Extensible: http://hl7.org/fhir/R4/valueset-document-classcodes.html
DocumentReference.securityLabel	Extensible: http://hl7.org/fhir/R4/valueset-security-labels.html Informative: Note that the latest value set has many more values than historically used in the eHx. This is what FHIR DocumentReference allows.
DocumentReference.content.attachment.contentType	Required: https://www.hl7.org/documentcenter/public/standards/vocabulary/vocabulary_tables/infrastructure/vocabulary/mediaType.html
DocumentReference.content.attachment.language	Preferred: http://www.hl7.org/fhir/valueset-languages.html Required: http://www.hl7.org/fhir/valueset-all-languages.html (not further constrained)
DocumentReference.content.format	Extensible: http://hl7.org/fhir/R4/valueset-formatcodes.html
DocumentReference.context.event	Example: http://hl7.org/fhir/R4/v3/ActCode/vs.html Informative: This specification does not constrain this attribute, as it is very specific to the type of document.
DocumentReference.facilityType	Extensible: http://hl7.org/fhir/R4/valueset-c80-facilitycodes.html
DocumentReference.practiceSetting	Extensible: http://hl7.org/fhir/R4/valueset-c80-practice-codes.html
DocumentManifest.type	Example: http://hl7.org/fhir/R4/v3/ActCode/vs.html (not further constrained)

Copyright© 2020-2021 All rights reserved.

The following indented elements apply to the PractitionerRole resource in the following paths: DocumentReference.author.ofType(PractitionerRole) DocumentReference.authenticator.ofType(PractitionerRole) DocumentManifest.author.ofType(PractitionerRole) DocumentManifest.recipient.ofType(PractitionerRole) Patient.generalPractitioner.ofType(PractitionerRole)	
PractitionerRole.code	Preferred: http://hl7.org/fhir/R4/valueset-practitioner-role.html
PractitionerRole.specialty	Preferred: http://hl7.org/fhir/R4/valueset-c80-practice-codes.html
The following indented elements apply to the Practitioner resource in the following paths: DocumentReference.author.ofType(Practitioner) DocumentReference.authenticator.ofType(Practitioner) DocumentManifest.author.ofType(Practitioner) DocumentManifest.recipient.ofType(Practitioner) PractitionerRole.practitioner Patient.generalPractitioner.ofType(Practitioner)	
Practitioner. photo.language	Not constrained – this field will be ignored
Practitioner.qualification. code	Example: http://www.hl7.org/fhir/v2/0360/2.7/index.html (not further constrained)
Practitioner. communication	Preferred: http://www.hl7.org/fhir/valueset-languages.html Required: http://www.hl7.org/fhir/valueset-all-languages.html (not further constrained)
The following indented elements apply to the Organization resource in the following paths: DocumentReference.author.ofType(Organization) DocumentReference.authenticator.ofType(Organization) DocumentManifest.author.ofType(Organization) DocumentManifest.recipient.ofType(Organization) DocumentReference.custodian PractitionerRole.organization PractitionerRole.practitioner.qualification.issuer	

Copyright© 2020-2021 All rights reserved.

Practitioner.qualification.issuer Location.managingOrganization HealthcareService.providedBy Device.owner DeviceDefinition.manufacturerReference DeviceDefinition.owner Patient.contact.organization Patient.generalPractitioner.ofType(Organization) Patient.managingOrganization Organization.repeat(partOf)	
Organization.type	Example: http://www.hl7.org/fhir/valueset-organization-type.html (not further constrained)
Organization.contact.purpose	Extensible: https://www.hl7.org/fhir/valueset-contactentity-type.html (not further constrained)
Organization.endpoint	Note: May want to consider whether this whole element if present should be constrained as it appears in the directory. That would take care of more than just coded values in it. I also would not keep it in this table but would have a separate requirement.
The following indented elements apply to the Location resource in the following paths: PractitionerRole.location HealthcareService.location HealthcareService.coverageArea Device.location Location.repeat(partOf)	
Location.operationalStatus	Preferred: http://hl7.org/fhir/R4/v2/0116/index.html (not further constrained)
Location.type	Extensible: http://hl7.org/fhir/R4/v3/ServiceDeliveryLocationRoleType/vs.html (not further constrained)
Location.physicalType	Example: http://hl7.org/fhir/R4/valueset-location-physical-type.html (not further constrained)

Copyright© 2020-2021 All rights reserved.

The following indented elements apply to the HealthcareService resource in the following paths: PractitionerRole.healthcareService	
HealthcareService. category	Example: http://hl7.org/fhir/R4/valueset-service-category.html (not further constrained)
HealthcareService.type	Example: http://hl7.org/fhir/R4/valueset-service-type.html (not further constrained)
HealthcareService. specialty	Preferred: http://hl7.org/fhir/R4/valueset-c80-practice-codes.html (not further constrained)
HealthcareService. serviceProvisionCode	Example: http://hl7.org/fhir/R4/valueset-service-provision-conditions.html (not further constrained)
HealthcareService. eligibility.code	None defined
HealthcareService. program	Example: http://hl7.org/fhir/R4/valueset-program.html (not further constrained)
HealthcareService. characteristic	None defined
HealthcareService. communication	Preferred: http://www.hl7.org/fhir/valueset-languages.html Required: http://www.hl7.org/fhir/valueset-all-languages.html (not further constrained)
HealthcareService. referralMethod	Example: http://hl7.org/fhir/R4/valueset-service-referral-method.html (not further constrained)
The following indented elements apply to the Device resource in the following paths: DocumentReference.author.ofType(Device) DocumentManifest.author.ofType(Device) Device.repeat(parent)	
Device.statusReason	Extensible: http://hl7.org/fhir/R4/valueset-device-status-reason.html (not further constrained)
Device.type	Example: http://hl7.org/fhir/R4/valueset-device-type.html (not further constrained)

Copyright© 2020-2021 All rights reserved.

Device.specialization.systemType	None defined
Device.version.type	None defined
Device.property.type	None defined
Device.property.valueCode	None defined
Device.safety	Example: http://hl7.org/fhir/R4/valueset-device-safety.html
The following indented elements apply to the DeviceDefinition resource in the following paths: Device.definition DeviceDefinition.repeat(parentDevice)	
DeviceDefinition.type	Example: http://hl7.org/fhir/R4/valueset-device-type.html (not further constrained)
DeviceDefinition.safety	Example: http://hl7.org/fhir/R4/valueset-device-safety.html (not further constrained)
DeviceDefinition.languageCode	Required: http://www.hl7.org/fhir/valueset-all-languages.html
DeviceDefinition.capability.type	None defined
DeviceDefinition.capability.description	None defined. Note: Likely an oversight – this should be a string, not coded.
DeviceDefinition.property.type	None defined
DeviceDefinition.property.valueCode	None defined
DeviceDefinition.material.substance	None defined
The following indented elements apply to the Patient resource in the following paths: DocumentReference.subject DocumentReference.author.ofType(Patient)	

Copyright© 2020-2021 All rights reserved.

DocumentReference.context.sourcePatientInfo DocumentManifest.subject DocumentManifest.author.ofType(Patient) DocumentManifest.recipient.ofType(Patient) Device.patient RelatedPerson.patient Patient.link.other.ofType(Patient)	
Patient.maritalStatus	Extensible: http://hl7.org/fhir/R4/valueset-marital-status.html (not further constrained)
Patient.photo.language	Not constrained – this field will be ignored
Patient. contact.relationship	Extensible: http://hl7.org/fhir/R4/valueset-patient-contactrelationship.html (not further constrained)
Patient.communication	Preferred: http://www.hl7.org/fhir/valueset-languages.html Required: http://www.hl7.org/fhir/valueset-all-languages.html (not further constrained)
The following indented elements apply to the RelatedPerson resource in the following paths: DocumentReference.author.ofType(RelatedPerson) DocumentManifest.author.ofType(RelatedPerson) DocumentManifest.recipient.ofType(RelatedPerson) Patient.link.other.ofType(RelatedPerson)	
RelatedPerson. relationship	Preferred: http://hl7.org/fhir/R4/valueset-relatedperson-relationshiptype.html (not further constrained)
RelatedPerson. photo.language	Not constrained – this field will be ignored
RelatedPerson. communication	Preferred: http://www.hl7.org/fhir/valueset-languages.html Required: http://www.hl7.org/fhir/valueset-all-languages.html (not further constrained)
TODO Encounter	
TODO EpisodeOfCare	

Copyright© 2020-2021 All rights reserved.

Note: The FHIR elements above are expressed using FHIRPath syntax: http://hl7.org/fhirpath/N1/ .

Informative: The above binding strengths are the same for the equivalent metadata elements in XDS (see section TBD).

4.12 Initiating Participant Processing Requirements: FHIR MHD Option

Note: In this section, an Initiating Participant that supports the FHIR MHD transaction option is referred to as a "FHIR MHD Initiating Participant", and a Responding Participant is similarly referred to as a "FHIR MHD Receiving Participant".

CONF-172: A FHIR MHD Initiating Participant SHALL implement the ITI-65 transaction as specified in [IHE MHD IG] Provide Document Bundle [ITI-65] and constrained in this section.

Informative: ITI-65 allows a document to be included directly in the Provide request as a Binary Resource or referenced via an absolute URL to where it is hosted on a server, which could be the sending system or a third party. The Binary flavor is like the SOAP Push mechanism, while the reference flavor has no SOAP counterpart. It could support cases like:

- A lighter push mechanism of just the metadata, where the server can choose to retrieve only the actual documents it wants.
- Pushing a reference to a third-party server, for example a consent server.

See an example ITI-65 request here: TBD.

CONF-173: A FHIR MHD Initiating Participant SHALL constrain DocumentReference.content.attachment.url to be one of the following:

- A URL that resolves to a Binary Resource included in the ITI-65 request.
- An absolute URL pointing to a document hosted at a server listed in the eHealth Exchange directory.

Informative: In the case of a FHIR document, there should be an option to pass it as a Bundle Resource rather than an encoded Binary Resource, making the message more readable. We have submitted this as a CP to IHE: see https://github.com/IHE/ITI.MHD/issues/34 .

4.12.1 Metadata Elements

Informative: Section [IHE MHD IG] 11.4.1.2.1 Bundle Resources gives clear requirements for the metadata resources (DocumentReference, List) across all variations of options. Follow the hyperlinks for the applicable StructureDefinitions. Note that in some cases, you have to

Copyright© 2020-2021 All rights reserved.

navigate to parent types. This section adds constraints to the MHD document metadata as expressed in the DocumentReference and List resources.

Informative: Note that the attribute optionality above applies to each specific resource that is submitted and is different than the ability to support a given attribute. For example, an Initiating Participant that does not support the Comprehensive Metadata Option may omit DocumentReference.category for a given document but still must support the ability to submit the category if one is available.

Some of the key metadata elements are further described and constrained here:

4.12.1.1 DocumentReference.context.sourcePatientInfo

CONF-174: If included, sourcePatientInfo SHOULD specify a minimum of demographics for the patient as known by the FHIR MHD Initiating Participant, including first name, last name, date of birth and gender.

Informative: This is an optional element, required if the Comprehensive Metadata Option is supported, required if known otherwise.

CONF-175: The business identifier in the sourcePatientInfo Patient Resource SHALL contain:

- type: “MR” for Medical Record Number
- system: Patient Identity Assigning Authority in the form of an OID URN
- value: An identifier in the above Assigning Authority domain

Informative: The business identifier in the sourcePatientInfo Patient Resource represents the community identifier of the subject of care (i.e., patient) of the document from the Initiating Participant’s Assigning Authority domain. See [IHE ITI TF-2] Appendix Z.9.1 Identifier Type for coding.

4.12.1.2 DocumentReference.subject

CONF-176: The business identifier in the subject Patient Resource, if included, SHALL contain:

- type: “MR” for Medical Record Number
- system: Patient Identity Assigning Authority in the form of an OID URN
- value: An identifier in the above Assigning Authority domain

Informative: The subject represents the subject of care of the document (i.e., patient), and is, if possible, hosted at the Receiving Participant’s server. This resource is obtained by the Initiating Participant through some verifiable means, primarily through use of the Patient Discovery

Copyright© 2020-2021 All rights reserved.

Specification. It is constrained to this format to be compatible with patient IDs obtained via XCPD. See [IHE ITI TF-2] Appendix Z.9.1 Identifier Type.

CONF-177: A FHIR MHD Initiating Participant SHALL populate the subject resource in all of the following cases:

- The Receiving Participant declares the eHx Patient Required Option.
- The Receiving Participant declares the Comprehensive Metadata Option.

CONF-178: A FHIR MHD Initiating Participant that omits the subject resource SHALL include context.sourcePatientInfo, and its demographics SHALL meet the conditions in [IHE ITI TF-2] 3.80.4.1.1.

Informative: The above requirement requires adequate demographics to ensure identification of the patient.

CONF-179: A FHIR MHD Initiating Participant, if populating the subject resource, SHOULD populate with the subject of care of the submission from the FHIR MHD Receiving Participant's Assigning Authority domain.

Informative: The reason the above is a SHOULD is to allow an initiator to still Push if the patient match fails yet the subject is required.

4.12.1.3 DocumentReference.content.attachment.hash

CONF-180: A FHIR MHD Initiating Participant SHALL populate the Hash with the hash of the document, computed following the SHA-1 algorithm.

4.12.1.4 DocumentReference.content.attachment.size

CONF-181: A FHIR MHD Initiating Participant SHALL populate the Size with the actual size (in bytes) of the document.

4.12.1.5 Submission Set List.subject

CONF-182: A FHIR MHD Initiating Participant SHOULD populate List.subject for the List resource representing the Submission Set with the subject of care of the submission from the FHIR MHD Receiving Participant's Assigning Authority domain.

CONF-183: The subject SHALL follow the same rules as defined for DocumentReference.subject in Section TBD and SHALL contain the same value if supplied.

Copyright© 2020-2021 All rights reserved.

4.12.1.6 Endpoint

Informative: An Endpoint resource (https://www.hl7.org/fhir/endpoint.html) may be included in a number of resources within the submission metadata: Organization, PractitionerRole, Location, and HealthcareService. There are two typical uses by clients:

- Searching a directory, e.g., of Organizations and Locations, for the base URLs for each organization and mechanism for contact. This is covered in the eHx Directory specification.
- Identifying an endpoint that has ultimate responsibility for a specific resource, if it is a different system or communications mechanism than that of the resource itself. This is covered in this section.

This second use is intended only to identify external systems, not the same FHIR server as the resource itself. So, for example, if you were to simply query a FHIR server for an Organization and its parent Organizations, the Organization resources you get back would not have to populate Organization.endpoint.

In a cross-community push, by contrast, Endpoint can become very important. Think of it as the FHIR version of a Health Data Locator, or Home Community ID, indicating the location/system that holds clinical information. For example, if an Initiating Participant submits a Patient resource with the following:

Patient.managingOrganization.endpoint.address = (server B)

this means that although there is a Patient resource included in the submission, the patient is actually managed by server B, and if the endpoint is a FHIR server, the Receiving Participant could use it to query for server B's resource for that patient, as well as other information about that patient, e.g., Encounters.

Alternatively, the patient could simply point to a resource at server B:

Patient.managingOrganization.reference = [server B]/Organization/123

CONF-184: When including an Endpoint resource anywhere in a submission, a FHIR MHD Initiating Participant SHALL ensure that the Endpoint.address is a URL to a FHIR server listed in the eHealth Exchange directory.

4.12.2 Submitting document relationships

Informative: Document-related information in FHIR MHD consists of a related and profiled set of resources: DocumentReference, Binary, List and Bundle. This information must be

Copyright© 2020-2021 All rights reserved.

immutable[1] once exchanged. Therefore, **Document-related resources are not updated by pushing new versions of resources using "update" or "patch".** Rather, updates are **new resources**, explicitly related to the resource(s) being updated.

CONF-185: A FHIR MHD Initiating Participant MAY submit any DocumentReference.relatesTo.code value. See http://hl7.org/fhir/R4/valueset-document-relationship-type.html.

Informative: This specification does not specify the conditions that trigger an Initiating Participant to submit a document relationship. Further, there should not be an expectation that the Receiving Participant will always accept the relationship—it may apply checks and processing before accepting, especially if the source system or author differs.

4.12.2.1 Reflecting the relationship in the DocumentReference

CONF-186: When relating to a document, a FHIR MHD Initiating Participant SHALL reference existing metadata according to Section 4.11.1 Submitting resource references.

CONF-187: When relating to a known DocumentReference resource, a FHIR MHD Initiating Participant SHALL populate DocumentReference.relatesTo.target.reference in the submission with the version-independent URL of the original DocumentReference.

CONF-188: If an attempt to relate to a DocumentReference fails due to a version clash (XDSRegistryDeprecatedDocumentError is returned), the FHIR MHD Initiating Participant SHOULD either abandon the attempt, or query to determine the latest applicable DocumentReference and submit a relationship to that instead. Note that the latest document may have relationships to appendices or transformations.

4.12.2.2 Reflecting the relationship in a CDA Document

CONF-189: When relating an HL7 CDA document to another CDA document, a FHIR MHD Initiating Participant SHALL follow the requirements in section 3.10.3, mapping codes as follows:

- replaces -> RPLC
- appends -> APND
- transforms -> XFRM

---

[1] Immutability of exchanged documents is a generally accepted concept, although it is not comprehensively addressed in FHIR. It is briefly addressed for FHIR documents ("A document is an immutable set of resources..." https://www.hl7.org/fhir/documents.html#3.3), but not for DocumentReference, Binary, or List in base FHIR or in MHD.

Copyright© 2020-2021 All rights reserved.

- signs: no document relationship needed

4.12.2.3 Reflecting the relationship in a FHIR Document

CONF-190: When relating an HL7 FHIR document to another FHIR document, a FHIR MHD Initiating Participant that has control over the generation of the new document SHALL populate Composition.relatesTo.code = DocumentReference.relatesTo.code and Composition.relatesTo.target = the prior document's Composition.

4.12.2.4 Workflow for Initiating Participants that can submit corrections

Within this section, the term "Correcting Initiating Participant" is defined as an Initiating Participant that supports the FHIR MHD transaction option and has the capability of submitting corrections to documents.

Informative: This section defines a conditional workflow for submitting content and subsequent corrections that handles the cases of version clashes as well as when a document was saved in some other way.

CONF-191: A Correcting Initiating Participant SHALL populate at least one DocumentReference.identifier in submissions.

CONF-192: A Correcting Initiating Participant SHALL process an ITI-65 transaction response as follows:

- If the Bundle.entry.response.status for a DocumentReference is 201, the Correcting Initiating Participant SHALL persist the version-independent URL (i.e., the URL without the _history path) from the Bundle.entry.response.location element for later potential corrections.
- If the Bundle.entry.response.status for a DocumentReference is 200, the Correcting Initiating Participant SHALL persist the submitted DocumentReference.identifier(s) for later potential corrections.

CONF-193: When updating a previously submitted document that was persisted as a DocumentReference, a Correcting Initiating Participant SHALL follow the requirements in section 4.12.2.1.

CONF-194: When updating a previously submitted document that was persisted in some other way than a DocumentReference, a Correcting Initiating Participant SHALL populate DocumentReference.relatesTo.target.identifier(s) in the submission with the value(s) originally submitted.

Copyright© 2020-2021 All rights reserved.

4.12.3 Submitting Folders and Submission Sets

Informative: Folders in FHIR are represented by a profiled List resource. Folder metadata is described in [IHE MHD] TBD.

Informative: A FHIR MHD Initiating Participant MAY submit folders, but Receiving Participants are not required to support them. Folders are not included because there has been no use case presented for them. This could change in the future.

Informative: Submission Sets in FHIR are represented by a profiled List resource. Submission set metadata is described in [IHE MHD] TBD. Note that Receiving Participants may or may not persist Submission Sets.

CONF-195: A FHIR MHD Initiating Participant MAY reference existing DocumentReferences in submitted Submission Sets, and these DocumentReferences MAY be for more than one patient.

Informative: The above requirement mirrors the behavior in XDS, which supports edge use cases like associating mother and child documents in one prenatal submission set. See [IHE ITI TF-3] 4.2.2.1.1 SS-DE HasMember.

CONF-196: A FHIR MHD Initiating Participant SHALL use only version-independent resource references when referencing Folder List resources as entries in Submission Set or Folder List resources.

4.12.4 Routing

Informative: An Initiating Participant wishing to route to additional recipients should ensure that the Receiving Participant recognizes all of them, for example, by determining the relationship to the Receiving Participant from the eHx directory. Any unknown or unreachable participant may cause the entire transaction to fail.

Informative: A FHIR MHD Initiating Participant passes the Home Community ID of the ultimate recipient in the appropriate fields as specified TBD.

CONF-197: A FHIR MHD Initiating Participant wishing to submit to a federated community SHALL populate exactly one Submission Set List.extension:intendedRecipient as follows:

- With a Resource type of Organization
- With exactly one Organization.identifier with .use of “official”
- Where that identifier is a Home Community ID found in the eHx directory

Copyright© 2020-2021 All rights reserved.

CONF-198: An Initiating Participant wishing to submit to a non-community organization or individual SHALL populate Submission Set List.extension:intendedRecipient with a reference identifying each ultimate recipient.

Informative: This mechanism isn't fully defined in MHD yet. We are addressing this in the IHE ITI MHD to a Federation work item (https://github.com/IHE/IT-Infrastructure/issues/142).

4.12.5 Submitting FHIR Documents

CONF-199: When submitting a FHIR document, a FHIR MHD Initiating Participant SHALL reference the document as either a Binary resource or as a Bundle resource containing the native FHIR document.

Informative: The MHD profile in section 3.65.4.1.2 only allows for Binary Resources to be included. But US Core (https://build.fhir.org/ig/HL7/US-Core-R4/StructureDefinition-us-core-documentreference.html) allows DocumentReference to point to a FHIR Document Bundle as well. As this results in a more readable message, we are adopting this broader definition and will submit a CP to update MHD.

Informative: FHIR documents are described here: https://www.hl7.org/fhir/documents.html. An example Discharge Summary is here: https://www.hl7.org/fhir/document-example-dischargesummary.json.html. This specification does not constrain documents to be expressed as JSON or XML.

Informative: The FHIR Documents page has one misleading statement: "FHIR documents are for documents that are authored and assembled in FHIR, while the document reference resource is for general references to pre-existing documents." The two are not mutually exclusive. A document reference may point to a fully assembled FHIR document.

Informative: In the context of a push, a FHIR document must be "an immutable set of resources with a fixed presentation that is authored and/or attested by humans, organizations and devices", for example, a document type Bundle after calling the Composition $document operation. This section (https://www.hl7.org/fhir/documents.html#content) defines the specific resources referenced within the document that must be included. Other referenced resources may be included or be resolvable only on a server. See the example document at https://www.hl7.org/fhir/document-example-dischargesummary.json.html. The net effect of this flexibility is that the sender can choose which, if any, of the resources included in a FHIR document to expose as independently referenceable FHIR resources.

Informative: Although the document and any included entries are an immutable snapshot in time, if the fullURL for an included entry in the document Bundle is a resolveable server URL,

Copyright© 2020-2021 All rights reserved.

that is a version-independent resource that the Receiving Participant can choose to cache and keep track of. Consider these possibilities:

1. Included resources are all UUIDs/OIDs: this is closest to pushing a CDA document, as it is fully self-contained.
2. Included resources, use URLs on the source server, the destination server, or a third party server. Example: system A pushes document to system B. In the document is a URL reference to a patient resource at system A. This allows system B to track the patient directly if needed.

4.13 Initiating Participant Processing Requirements: FHIR Resource Option

Informative: Pushing individual resources across communities is not a typical use case. Usually, rather than simple acceptance and persistence of the resource by the server, such a push will trigger a workflow. This section describes and constrains the acceptable cases.

CONF-200: An Initiating Participant utilizing the FHIR Resource Option SHALL be constrained to the cases defined in this section, or in another eHealth Exchange profile.

CONF-201: An Initiating Participant utilizing the FHIR Resource Option SHALL be constrained to the following mechanisms for submitting resources:

- create (POST: see https://www.hl7.org/fhir/http.html#create)
- update (PUT: see https://www.hl7.org/fhir/http.html#update)
- batch/transaction (POST /Bundle with POST or PUT entries: see https://www.hl7.org/fhir/http.html#transaction)

Informative: Note that the **update as create** (https://www.hl7.org/fhir/http.html#upsert) and **patch** (https://www.hl7.org/fhir/http.html#patch) interactions are not currently supported for pushing content.

CONF-202: An Initiating Participant utilizing the FHIR Resource Option SHOULD check the Receiving Participant's Capability Statement to determine whether it supports the specific use case before presenting to an end-user the ability to execute the use case to this specific system.

4.13.1 Use Case: Patient Correction

Informative: The Patient Correction use case is owned and defined by the HL7 Patient Empowerment Workgroup.

CONF-203: An Initiating Participant utilizing the FHIR Resource Option for the Patient Correction use case SHALL implement the requirements for clients in [HL7 Patient Correction].

Copyright© 2020-2021 All rights reserved.

Informative: This use case requires no separate Provenance resource to be submitted.

4.14 Initiating Participant Processing Requirements: FHIR Messaging Option

This section is reserved.

4.15 FHIR Processing Requirements: Provenance

CONF-204: The Initiating Participant SHALL populate provenance information for FHIR resources in submissions according to this specification, unless overridden by another eHx Provenance specification.

Informative: The purpose of this requirement is to allow future content-specific profiles to further relax or constrain provenance as defined in this document to meet use cases.

Informative: Provenance records the key information about the "5 w's" (i.e., who, what, etc.) involved in the creation or modification of data. See the Provenance resource: https://www.hl7.org/fhir/provenance.html. In general, the rule is that if the information is already accounted for in the resource, then use that, but, if it isn't, use a separate Provenance Resource that points to the resource in question.

In the context of a cross-community submission, there is always a need to convey provenance to the receiver.

In general we follow the guidance of [HL7 Basic Provenance] which covers this use case at a high level, focusing on minimal requirements for author, author organization and timestamp, but with the following exception. In the case of FHIR resources, it adopts the US Core Provenance profile, which requires Provenance resources for a number of resources, including DocumentReference. However, it does not cover the submission case in detail, and does not reference ITI-65 or any IHE profiles. We believe that when the IHE ITI-65 transaction is used, the overall submission provides more than adequate provenance to meet basic requirements without needing to include explicit Provenance resources in the submission. We plan to submit change requests to this effect.

Informative: Provenance needs for Patient Generated Health Data (PGHD) have not been fully analyzed in the industry. Numerous groups (e.g. ONC, HL7 Patient Empowerment) have looked into the topic, but to date there are no governing profiles for these use cases. We will assist with this, and until nailed down will adopt minimal requirements. We will look to cover the following cases:

Copyright© 2020-2021 All rights reserved.

- If the source is the patient acting through a sending system that is a member of the eHx (e.g. a patient portal via browser or consumer app)
- If the source is a patient acting through a consumer app
- If the source is a device owned and operated by the patient
- If the source is a device owned and operated by the provider and in the custody of the patient

4.15.1 Reflecting Provenance for a Submission Set

CONF-205: When using the ITI-65 transaction, an Initiating Participant SHALL populate provenance information within the Submission Set List resource as follows:

- If the source is a provider, List.source SHALL be of type PractitionerRole and PractitionerRole.organization.name and PractitionerRole.practitioner.name SHALL be present.
- If the source is a device in a care setting, List.source SHALL be of type Device and Device.owner.name and Device.identifier SHALL be present.
- If the source is the patient acting through a sending system (e.g., a patient portal via browser or consumer app), List.source SHALL be of type Patient and Patient.name SHALL be present.

4.15.2 Reflecting Provenance for submitted C-CDA Documents

CONF-206: When submitting HL7 C-CDA documents, an Initiating Participant SHALL follow the Provenance requirements in section 3.14.

Informative: The above requirement is irrespective of the transport mechanism.

4.15.3 Reflecting Provenance for submitted FHIR Documents

CONF-207: When submitting HL7 FHIR documents, an Initiating Participant SHALL follow the Provenance requirements in this section.

Informative: The above requirement is irrespective of the transport mechanism.

4.15.3.1 Updating a document: Appending or replacing from the same source

Informative: When updating an HL7 FHIR document, the association between the original and new document is considered an important part of provenance and is fully specified in other sections of this specification. The following cases for updating a document do not require any additional requirements for conveying provenance:

- Submitting an appendix to a document.

Copyright© 2020-2021 All rights reserved.

- Submitting a replacement document from the same source as the prior document (same source system, community and authors).

4.15.3.2 Updating a document: Replacing from a different source

CONF-208: If an Initiating Participant is submitting a replacement to an HL7 FHIR document originally from another source, defined as a different community or author, and wishes to add or modify only individual entries, it SHALL:

- Use the appropriate document-level template from [HL7 Data Provenance] depending on the kind of author, e.g., Provider Generated Document With Provenance.
- Replace the document-level author, and use the original author for those sections or entries which are not being changed.
- Use the appropriate entry-level template from [HL7 Data Provenance] for any modified entry, e.g., Observation Generated by Provider.

CONF-209: If an Initiating Participant is submitting a replacement to an HL7 FHIR document originally from another source, defined as a different community or author, and wishes to reconcile entire sections, it SHALL:

- Use the appropriate document-level template from [HL7 Data Provenance] depending on the kind of author, e.g., Provider Generated Document With Provenance.
- Replace the document-level author, and use the original author for those sections which are not being changed.
- Use the appropriate entry-level template from [HL7 Data Provenance] for any modified entry, e.g., Observation Generated by Provider.
- Conform to the content requirements in [IHE RECON] 6.3.1.D Reconciliation Content, including a Reconciliation Act in each section that has been reconciled.

Informative: For example: a patient's PCP submits a replacement document for one authored by another physician because it has incorrect information about the patient. We cover two such cases: selective correction and section-level reconciliation. Note that cross-author updates may be treated differently from ordinary submissions by the Receiving Participant, for example, additional security checks may be performed. See the Security Considerations section 5 for details.

4.15.4 Reflecting Provenance in FHIR Resources

CONF-210: The Initiating Participant SHALL populate provenance information in all submitted FHIR resources according to the rules in [HL7 Basic Provenance], constrained and clarified as follows:

Copyright© 2020-2021 All rights reserved.

- 3.1 Basic Provenance Practices:
  - CONF:1000. Informative: while this mentions care providing organizations as authors, we clarify that document content can be authored by other types of entities, e.g., patients, payers, labs, etc.
  - CONF:1004 and CONF:1005. Informative: Note that these require authorship to be captured for changed data, so that it is accurate when submitted to another system. While this behavior is upstream from the actual submission, we believe this should be standard practice already.
- 3.3 FHIR:
  - All referenced requirements.
  - **CONF-211**: The Initiating Participant SHALL include Provenance resources for each resource submitted using an individual PUT or POST, using the X-Provenance header as specified in [HL7 FHIR R4] Provenance section 6.3.4.2.1: https://www.hl7.org/fhir/provenance.html#header.
  - **CONF-212**: The Initiating Participant SHALL include Provenance resources for each resource submitted within a Bundle, by including those resources in the same Bundle.

4.16 Receiving Participant Common FHIR Requirements

CONF-213: FHIR Receiving Participants SHALL follow the requirements and guidance in [IHE ITI TF-2] Appendix Z for the "server actor", as well as any general requirements and guidance, e.g., Section Z.9 FHIR Data Types.

4.16.1 Creating resources

Informative: The "create" interaction is an HTTP POST command as described in [FHIR R4] 3.1.0.8 create (see https://www.hl7.org/fhir/http.html#create). When a Receiving Participant receives content in the form of POSTed FHIR resources, whether they are POSTed individually or included as POSTs in a transaction Bundle, the Participant has choices in how to process each resource. See [FHIR R4] 3.1.0.13 Transactional Integrity (https://www.hl7.org/fhir/http.html#transactional-integrity) and 7.14 Variations between Submitted data and Retrieved data (http://hl7.org/fhir/R4/updates.html#7.14). In a cross-community setting, these variations are more likely.

For example, any of the following could be considered successful processing:

- Persisting a submitted resource as is
- Adding additional data to the persisted resource (See 7.14.4 Generated and inferred data https://www.hl7.org/fhir/updates.html#7.14.4)

Copyright© 2020-2021 All rights reserved.

- Persisting the information from a resource in some other way, for example: adding a record of a COVID-positive case in a public health database
- Ignoring data or resources not considered relevant
- Merging or linking an external resource with an internal one, for example: a Patient resource from a third party is passed in, and it can be successfully matched with a resource at the Receiving Participant
- **De-containing** a contained resource (special case of the above). This means that a resource is submitted as a contained resource, but the receiver has additional knowledge to either match the resource to an existing one or to create a new resource. In this case, the resulting parent resource will have the reference to the contained resource updated to point to the separate resource.[2]

The following requirements apply to the create response for each resource, whether returned in the HTTP response or in Bundle.entry.response.

CONF-214: If the FHIR Receiving Participant creates a resource from the POSTed resource exactly:

- it SHALL return a status of 201 to indicate the Resource has been created.
- it SHALL populate Location with the version-specific URL (relative or absolute) of the resulting resource.
- it SHALL populate ETag with the resource version, if versioning is supported.

CONF-215: If the FHIR Receiving Participant creates a resource from the POSTed resource while modifying some information, or matches a POSTed resource with an existing resource, potentially merging in the submitted information:

- it SHALL return a status of 201 which will appear to the client as though the Resource has been created.
- it SHOULD return a copy of the newly created resource in the body of the response.

---

[2] See: https://chat.fhir.org/#narrow/stream/179166-implementers/topic/de-containment/near/224046501 . Some key quotes: Grahame Grieve: “The grounds on which this would be possible are that the server has access to business logic - specifically, statements about uniqueness that can be verified against some repository - that the client does not.” Lloyd McKenzie: “If a server says "this thing that you thought wasn't identifiable and had no independent existence, I can identify and choose to treat as independent", that's legitimate, though tricky. It may make subsequent updates from that client harder or impossible.”

Copyright© 2020-2021 All rights reserved.

- it SHALL populate Location with the version-specific URL (relative or absolute) of the resulting resource.
- it SHALL populate ETag with the resource version, if versioning is supported.

CONF-216: If the FHIR Receiving Participant persists a POSTed resource in some other way than as a single FHIR resource, or ignores the resource as not necessary or relevant for successful processing, such that subsequent GETs would not be able to return a resource,

- it SHALL return a status of 201 which will appear to the client as though the Resource has been created.
- it SHALL populate Location with a URL (relative or absolute) that corresponds to the POSTed resource.
- it MAY populate ETag with a resource version.
- if the information from the resource was persisted in any way, it SHOULD persist the mapping between the returned URL and the persisted data so that subsequent corrections may be accepted. See section 4.16.2.

CONF-217: If the FHIR Receiving Participant processes a POSTed resource within a Batch/Transaction Bundle by "de-containing" a contained resource into a separately managed resource, it SHALL NOT return the additional resource as an entry in the Bundle response.

Informative: This maintains the required one-to-one relationship between POSTed entries and responses as specified in [FHIR R4] section 3.1.0.11.3 Batch/Transaction Response (https://www.hl7.org/fhir/http.html#transaction-response). The client will still be able to locate the newly created resources through references in the parent resource.

Informative: See [FHIR R4] section 2.5.0.3 Exchanging Extensions (https://www.hl7.org/fhir/extensibility.html#exchange) for additional guidance when creating resources containing extensions.

4.16.2 Updating resources

CONF-218: A FHIR Receiving Participant that had persisted information from a POSTed resource in some other way than an addressable resource SHOULD accept updates against the resource URL it returned in the create response, updating the underlying data as appropriate.

CONF-219: A FHIR Receiving Participant that persists clinical content as FHIR Resources and allows cross-community updates SHALL be able to handle conflicting updates. Two mechanisms are specified in this section as well as guidance for when they cannot be used. Mechanisms MAY differ for different resource types or use cases.

Copyright© 2020-2021 All rights reserved.

CONF-220: A FHIR Receiving Participant that persists clinical content as FHIR Resources and allows cross-community updates SHALL reject a PUT method on an individual resource or a PUT method within a transaction/batch Bundle if it detects the update cannot be done (e.g., due to server side pessimistic locking), returning a 409 Conflict status code.

4.16.2.1 Handling conflicting updates through If-Match

CONF-221: A FHIR Receiving Participant that handles conflicting updates through If-Match SHALL support resource record versions—see [FHIR R4] section 2.26.3.7.1 Record Versions vs Business Versions vs FHIR Versions, https://www.hl7.org/fhir/resource.html#versions. FHIR Receiving Participants MAY retain resource history.

CONF-222: A FHIR Receiving Participant that handles conflicting updates through versions SHALL reject a PUT method on an individual resource or a PUT method within a transaction/batch Bundle that does not include the If-Match header, returning a 400 Client Error status code.

CONF-223: A FHIR Receiving Participant that handles conflicting updates through versions SHALL reject a PUT method on an individual resource or a PUT method within a transaction/batch Bundle If the version ID given in the If-Match header does not match the current version of the resource, returning a 412 Precondition Failed status code.

Informative: The If-Match mechanism as specified is in conflict with the underlying RFC 7232 (see https://jira.hl7.org/browse/FHIR-31925). Discussion with FHIR subject matter experts and implementers reflects that systems are proceeding with the behavior as specified by FHIR. We anticipate this wil be clarified in a future version of FHIR.

4.16.2.2 Handling conflicting updates through If-Unmodified-Since

Informative: [FHIR R4] section 3.1.0.1.12 Support for Versions (https://www.hl7.org/fhir/http.html#versions) states “Servers that do not support versioning SHALL ensure that Resource.meta.versionId is not present on resources they return, and SHALL update the value of Resource.meta.lastUpdated correctly.” Because servers must supply this value, it may be used by clients in updates to prevent collisions. However, note that this mechanism is not mentioned in [FHIR R4], and is not supported within a transaction Bundle. See https://jira.hl7.org/browse/FHIR-31927.

CONF-224: A FHIR Receiving Participant that handles conflicting updates through If-Unmodified-Since SHALL follow the requirements in section 3.4 of RFC 7232 (see https://tools.ietf.org/html/rfc7232#section-3.4).

Copyright© 2020-2021 All rights reserved.

4.16.2.3 Handling conflicting updates through other means

Informative: A FHIR Receiving Participant may detect and handle conflicting updates through other means. There are no specific normative requirements, but some suggestions follow:

- Consider locking a resource during update. This will only prevent race conditions, not content collisions, however.
- If the business version is used, check for two updates that claim the same new business version.
- If some fields can be considered immutable after creation, or changeable only by certain clients (e.g. those with local admin privileges), any changes beyond this can be rejected (with a 412 Precondition Failed) or silently ignored. For example, in the eHx directory the Administrative Status of an Organization (Terminated, Active, Testing) cannot be changed by a participant using a PUT or POST; it can only be changed by eHx staff.
- If the nature of the resource is that content may only be added to it, not removed, then conflicting updates may be gracefully handled by simply ignoring any removed content. This technique will work for the MHD Folder use of the List resource, and may work for the Patient Correction use of the Task resource (specifically Task.note and Task.output).

4.16.2.4 Update MHD-related resources only using MHD

CONF-225: A FHIR Receiving Participant SHALL reject attempts to update the following document-related resources by means of a PUT method on an individual resource or a PUT method within a transaction/batch Bundle, returning a 405 Method Not allowed response code:

- A DocumentReference
- A Binary or a Bundle referenced by a DocumentReference
- A List profiled as an MHD Submission Set

CONF-226: A FHIR Receiving Participant SHALL reject attempts to update a List profiled as an MHD Folder by means of a PUT method on an individual resource or a PUT method within a transaction/batch Bundle that is not profiled as an MHD ITI-65 transaction, returning a 405 Method Not allowed response code.

Informative: Documents and related metadata may only be updated using the FHIR MHD option. See section 4.17.

4.16.3 Handling resource references

Informative: See section 4.11.5, “Submitting resource references” for the allowable variations in references.

Copyright© 2020-2021 All rights reserved.

For resolving references in Bundles, see http://hl7.org/fhir/bundle.html#references as well as the proposed update to this text: https://jira.hl7.org/browse/FHIR-29271?focusedCommentId=183020&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-183020.

Note that a fullUrl for a resource that is to be POSTed is ONLY needed in order to resolve in-bundle references from other resources, and it is to be ignored otherwise. See [FHIR R4] section 3.1.0.11.2 Transaction Processing Rules (https://www.hl7.org/fhir/http.html#trules).

CONF-227: FHIR Receiving Participants SHALL validate any local URL and business identifier resource references when initially processing the FHIR Push Transaction, incorporating any errors in the response.

CONF-228: FHIR Receiving Participants SHALL attempt to fetch any external resource references that will be needed when initially processing the FHIR Push Transaction, incorporating any errors in the response.

CONF-229: FHIR Receiving Participants SHALL have a policy that defines which external resource references they will need to retrieve and which ones are critical, e.g., for processing, legal, archival purposes, etc.

CONF-230: If a FHIR Receiving Participant cannot resolve a resource referenced by the submitted resource, it MAY reject the resource entirely, returning a 422 Unprocessable Entity response code. See [FHIR R4] 3.1.0.8 create: https://www.hl7.org/fhir/http.html#create

Informative: The FHIR MHD option has more specific error handling for missing resources. See section TBD.

CONF-231: FHIR Receiving Participants MAY store external resource references from FHIR Push Transactions for future use.

Informative: The above requirements ensure any problems with references are caught and returned, and preserve information at the time of submission in case necessary information changes or is removed.

These support a robust workflow for building and utilizing a network of fine-grained resource references for future clinical uses. For example:

- Receiving Participant receives FHIR Push Transaction
    - fetches and saves the content of external resources
    - saves the absolute URLs of external resources

Copyright© 2020-2021 All rights reserved.

  - makes initial use of the information
- Later, Receiving Participant makes additional use of the information
  - first attempts to retrieve external references by resolving absolute URLs, following redirects if necessary (to obtain the latest information)
  - next, if any information cannot be retrieved, makes use of saved information

4.17 Receiving Participant Processing Requirements: FHIR MHD Option

Note: In this section, *a Responding Participant is referred to as a* "FHIR MHD Receiving Participant*", and an Initiating* Participant that supports the FHIR MHD *transaction option is similarly referred to as a "FHIR MHD Initiating Participant".*

CONF-232: A FHIR MHD Receiving Participant SHALL implement the ITI-65 transaction as specified in [IHE MHD IG] Provide Document Bundle [ITI-65] and constrained in this section.

4.17.1 Overall Processing

CONF-233: The FHIR MHD Receiving Participant SHALL process the entire submission, including any federated submissions, before returning, with the following exceptions:

If a document needs to be queued for manual matching to a patient, the FHIR MHD Receiving Participant SHALL return a DocumentQueued warning code for each document so queued.

- If full receipt by an intendedRecipient is not possible. See section 4.17.5.

Informative: The base IHE transactions require full processing of the submission before returning. However, there is an XDR warning code, DocumentQueued, that appears to permit an exception to this expectation. We clarify that here.

CONF-234: If multiple exceptional conditions exist, the FHIR MHD Receiving Participant SHALL detect and include each one in the response.

CONF-235: If the FHIR MHD Receiving Participant determines that submitted document(s) and/or metadata will be persisted in any way, it SHALL validate and update select received metadata, as detailed below, prior to persisting anything.

CONF-236: If the FHIR MHD Receiving Participant determines that submitted document(s) and/or metadata will be persisted as FHIR resources, it SHALL:

1. Make any persisted documents (Binary resources) available for retrieval via the Retrieve Documents transaction.

Copyright© 2020-2021 All rights reserved.

2. Make any persisted metadata (DocumentReference and List resources) available for retrieval via the Query for Documents transaction.

CONF-237: If the FHIR MHD Receiving Participant determines that submitted document(s) and/or metadata will be persisted in any way, if any error is encountered during processing of this request, the FHIR MHD Receiving Participant SHALL revert any changes to its data store for metadata. The FHIR MHD Receiving Participant SHALL then terminate processing the transaction and return an error response.

CONF-238: If the FHIR MHD Receiving Participant determines that submitted document(s) and/or metadata will be persisted in any way, if any error is encountered during processing of this request, the FHIR MHD Receiving Participant SHOULD revert any changes to its data store for document content. The FHIR MHD Receiving Participant SHALL then terminate processing the transaction and return an error response.

Informative: The above requirement is needed to retain compatibility with a potential XDS infrastructure behind the interface, which allows for the document to be persisted even if the metadata is rejected.

4.17.2 Creating resources

CONF-239: The FHIR MHD Receiving Participant SHALL set the status of all List or DocumentReference resources it chooses to persist to “current”.

CONF-240: If the FHIR MHD Receiving Participant determines that List or DocumentReference resources will be persisted, it MAY save extra metadata attributes that were included:

- Extra metadata attributes are defined as extensions on the List or DocumentReference resources.
- If the FHIR MHD Receiving Participant saves the extra metadata attributes, it SHALL include that metadata in later query responses.
- If there are extra metadata attributes in the Submission and the FHIR MHD Receiving Participant does not save them, the FHIR MHD Receiving Participant SHALL include the XDSExtraMetadataNotSaved warning in its response.
- See [FHIR R4] section 2.5.0.3 Exchanging Extensions (https://www.hl7.org/fhir/extensibility.html#exchange) for additional requirements and guidance.

Informative: As explained in section 4.17.1, Creating resources, FHIR Receiving Participants may choose to persist submitted resources in some other way than as the submitted resource. This

Copyright© 2020-2021 All rights reserved.

typically means there is no way to access or correct the resource. However, for documents, we define a mechanism to do so below and in section 4.17.7.

CONF-241: If the FHIR MHD Receiving Participant persists a submitted document in some other way than as FHIR resources, it SHOULD persist any supplied DocumentReference resource business identifiers as well, in order to correlate any future updates to this information.

4.17.3 Handling resource references

CONF-242: If a DocumentReference.content.attachment.url contains a URL that does not resolve within the submission Bundle, the FHIR MHD Receiving Participant SHALL attempt to fetch the document when initially processing the FHIR Push Transaction, incorporating any errors in the response.

CONF-243: The FHIR MHD Receiving Participant SHALL return an XDSMissingDocument error if it cannot resolve a DocumentReference.content.attachment.url reference.

CONF-244: The FHIR MHD Receiving Participant SHALL return an XDSMissingDocumentMetadata error if a Binary or FHIR document Bundle is included in a submission without a corresponding DocumentReference.

4.17.4 Patient Matching

CONF-245: If a DocumentReference.subject is not included, the FHIR MHD Receiving Participant SHOULD attempt to determine the local patient to add the document to by matching the demographics in DocumentReference.context.sourcePatientInfo.

CONF-246: If a DocumentReference.subject is included but is unrecognized, the FHIR MHD Receiving Participant SHOULD attempt to determine the local patient to add the document to by matching the demographics in DocumentReference.context.sourcePatientInfo.

CONF-247: If no local patient for the document can be identified, the FHIR MHD Receiving Participant MAY return an XDSUnknownPatientId error code.

Informative: We are allowing receivers to be forgiving here—they may not need the idea of a patient ID at all.

4.17.5 Routing

Informative: FHIR MHD Receiving Participants that support routing to any federated communities, sub-organizations or persons will work with eHx staff or update the eHx Directory directly to ensure each potential recipient is represented appropriately in the directory, or is included in an external directory available to participants.

Copyright© 2020-2021 All rights reserved.

Informative: Submission Set List.extension:ihe-intendedRecipient processing is fully defined in [IHE MHD IG] 2:3.65.4.1.3 Expected Actions.

4.17.6 Handling documents

CONF-248: For each DocumentReference, the FHIR MHD Receiving Participant SHALL:

- Verify the hash element, which must be present. If a received hash value differs from the calculated hash of the received document, the FHIR MHD Receiving Participant SHALL return an XDSRepositoryMetadataError error.
- Verify the size element, which must be present. If a received size value differs from the octet count of the received document, the FHIR MHD Receiving Participant SHALL return an XDSRepositoryMetadataError error.

CONF-249: If DocumentReference.identifier(s) match an existing DocumentReference and the size or hash attributes differ, the FHIR MHD Receiving Participant SHALL return an XDSNonIdenticalSize or XDSNonIdenticalHash error as appropriate.

CONF-250: If DocumentReference.identifier(s) match an existing DocumentReference and the size or hash attributes match, the FHIR MHD Receiving Participant SHALL NOT return an error.

Informative: The multiple DocumentReferences mentioned in the above requirement may reference the same instance of the document or may reference different (identical) copies of the document (for example, stored on different Document Repositories). This is needed to retain compatibility with a potential XDS infrastructure behind the interface.

CONF-251: The FHIR MHD Receiving Participant SHALL return an XDSRegistryMetadataError or XDSRepositoryMetadataError error if a DocumentReference.context.period.start is later than DocumentReference.context.period.end.

CONF-252: The FHIR MHD Receiving Participant MAY return an InvalidDocumentContent error if the document content does not match the DocumentReference.

4.17.7 Handling document relationships

Informative: Document-related information in FHIR MHD consists of a related and profiled set of resources: DocumentReference, Binary, List and Bundle. This information must be immutable once exchanged, to preserve referential integrity. Therefore, **Document-related resources are not updated by pushing new versions of resources using "update" or "patch".** Rather, updates are **new resources**, explicitly related to the resource(s) being updated. Note that this restriction does not apply to additional resources referenced by these resources. For example, a Patient referenced by a document may continue to have additional addresses, alternate names, etc. added.

Copyright© 2020-2021 All rights reserved.

Informative: A FHIR MHD Receiving Participant is required by [IHE MHD IG] to support document relationships (DocumentReference.relatesTo). See http://build.fhir.org/ig/IHE/ITI.MHD/StructureDefinition-IHE.MHD.Minimal.DocumentReference.html.

CONF-253: A FHIR MHD Receiving Participant SHALL accept all DocumentReference.relatesTo.code values. See http://hl7.org/fhir/R4/valueset-document-relationship-type.html.

Informative: The DocumentReference.relatesTo.target reference will either be a URL to a DocumentReference, or a business identifier. It will be a business identifier in the case where an Initiating Participant is referencing a document that was not persisted with a DocumentReference. Both cases are covered below.

CONF-254: If a FHIR MHD Receiving Participant receives a DocumentReference with .relatesTo populated and the DocumentReference.relatesTo.target reference does not resolve within the Submission Bundle, the FHIR MHD Receiving Participant SHALL attempt to fetch the target DocumentReference or other data to be updated when initially processing the FHIR Push Transaction, incorporating any errors in the response.

CONF-255: If a FHIR MHD Receiving Participant cannot resolve a DocumentReference.relatesTo.target reference, the FHIR MHD Receiving Participant SHALL return an UnresolvedReferenceException error.

CONF-256: If a FHIR MHD Receiving Participant receives a DocumentReference.relatesTo.target that is an existing DocumentReference, the FHIR MHD Receiving Participant SHALL verify that both DocumentReferences refer to the same subject, and if not, return an XDSPatientIdDoesNotMatch error.

CONF-257: If a FHIR MHD Receiving Participant receives a DocumentReference.relatesTo.target that is information stored in some other way than a DocumentReference, the FHIR MHD Receiving Participant SHALL verify that the patient to be updated is the same as that in the new DocumentReference, and if not, return an XDSPatientIdDoesNotMatch error.

CONF-258: If a FHIR MHD Receiving Participant receives a DocumentReference.relatesTo.target that is an existing DocumentReference, the FHIR MHD Receiving Participant SHALL verify that the target DocumentReference.status = “current”; otherwise, it shall return the error XDSRegistryDeprecatedDocumentError. This ensures that only the most recent version of a document can be replaced, etc.

Copyright© 2020-2021 All rights reserved.

CONF-259: If a FHIR MHD Receiving Participant receives a DocumentReference.relatesTo.target that is information stored in some other way than a DocumentReference, the FHIR MHD Receiving Participant SHALL, if possible, verify that the target information stored is current; otherwise, it shall return the error XDSRegistryDeprecatedDocumentError. This ensures that only the most recent version of a document can be replaced, etc.

CONF-260: If a FHIR MHD Receiving Participant receives a DocumentReference.relatesTo.code value of "replaces" and the target is a DocumentReference in the same submission, the FHIR MHD Receiving Participant SHALL return the error UnresolvedReferenceException.

CONF-261: If a FHIR MHD Receiving Participant receives a DocumentReference.relatesTo.code value of "replaces" and the target is an existing DocumentReference, the FHIR MHD Receiving Participant SHALL ensure replacement semantics are followed by marking the target DocumentReference.status as "superseded", as well as any transformations or appendices of it.

CONF-262: If a FHIR MHD Receiving Participant receives a DocumentReference.relatesTo.code value of "replaces" and the target is information stored in some other way than a DocumentReference, the FHIR MHD Receiving Participant SHALL, if possible, ensure replacement semantics are followed as appropriate, for example: remove or mark the information as superseded.

Informative: One example of the above requirement would be Public Health Reporting, where the only thing persisted was a document identifier and a record of a condition. In this case, the Receiver would be required to update that record accordingly.

Informative: For the case of cross-author updates, see the Security Considerations section 5 for additional considerations.

4.17.8 Handling Folders and Submission Sets

Informative: the processing requirements for Folders are defined in [IHE MHD] Expected Actions (http://build.fhir.org/ig/IHE/ITI.MHD/branches/master/ITI-65.html#expected-actions).

CONF-263: If a submission contains a Folder List resource and the FHIR MHD Receiving Participant does not support Folders, it SHALL ignore the Folder, continuing processing of the transaction and return a "PartialFolderContentNotProcessed" warning.

CONF-264: If a FHIR MHD Receiving Participant persists Folders, it SHALL allow updating a Folder List resource via a PUT method for the resource in an ITI-65 request Bundle.

CONF-265: When processing a create or update (i.e., POST or PUT) of a Folder List resource, the FHIR MHD Receiving Participant SHALL verify that the subject of each List.entry is the same and

Copyright© 2020-2021 All rights reserved.

that it matches the List.subject if present and, if not, return an XDSPatientIdDoesNotMatch error.

CONF-266: When processing a Submission Set List resource, the FHIR MHD Receiving Participant SHALL verify that the subject of each List.entry that is a **new** DocumentReference is the same and that it matches the List.subject if present and, if not, return an XDSPatientIdDoesNotMatch error.

CONF-267: When processing a Submission Set List resource, the FHIR MHD Receiving Participant SHALL allow List.entry references to **existing** DocumentReference resources that are for a different patient.

Informative: The above two requirements about Submission Set mirror the behavior in XDS, which supports edge use cases like associating mother and child documents in one prenatal submission set. See [IHE ITI TF-3] 4.2.2.1.1 SS-DE HasMember.

CONF-268: When processing a create (i.e., POST) of a Submission Set or Folder List resource, the FHIR MHD Receiving Participant SHALL verify that each List.entry that is a DocumentReference has DocumentReference.status = “current”; otherwise, it shall return the error XDSRegistryDeprecatedDocumentError. This ensures that only the most recent version of a document can be added to a list.

CONF-269: When processing an update (i.e., PUT) of a persisted Folder List resource, the FHIR MHD Receiving Participant SHALL verify that each **new** List.entry that is a DocumentReference has DocumentReference.status = “current”; otherwise, it shall return the error XDSRegistryDeprecatedDocumentError.

Informative: The above requirement follows the pattern in XDS. It ensures that only the most recent version of a document can be added to a folder but also allows superseded documents to remain in the folder. See [IHE ITI TF-3] 4.2.2.2.3 RPLC.

CONF-270: When processing an update (i.e., PUT) of a persisted Folder List resource, if any entries in the Folder have been removed, the FHIR MHD Receiving Participant SHALL ignore those items to be removed, leaving the entries in the persisted Folder.

Informative: XDS only permits adding items to Folders, not removing them. We are retaining that restriction in FHIR so that existing Folder relationships are immutable.

CONF-271: When processing an update (i.e., PUT) of a persisted Folder List resource, if the FHIR MHD Receiving Participant persists historical versions, it SHALL change the List.status of the prior version to “retired”.

Copyright© 2020-2021 All rights reserved.

CONF-272: When processing a create (i.e. POST) of a Submission Set or Folder List resource, the FHIR MHD Receiving Participant SHALL verify that the List.identifer(s) do not match those of a previously persisted List; otherwise, it shall return the error XDSDuplicateUniqueIdInRegistry.

4.17.9 Additional Exception Checking

CONF-273: The FHIR MHD Receiving Participant SHOULD return an XDSRegistryMetadataError or XDSRepositoryMetadataError code with a severity of Warning if a coded value is submitted that is not within the defined value sets in Table 2 Value sets for FHIR resources.

Informative: see http://sequoiatechwg.editme.com/Topic-1537135046658 and https://healthcaresecprivacy.blogspot.com/2011/11/xdsxca-testing-of-vocabulary.html.

CONF-274: The FHIR MHD Receiving Participant SHALL return an XDSRegistryDuplicateUniqueIdInMessage or XDSRepositoryDuplicateUniqueIdInMessage error if a resource business identifier value was used in more than one resource within the submission.

CONF-275: The FHIR MHD Receiving Participant MAY return the following error codes as defined in [IHE ITI TF-3] Table 4.2.4.1-2: XDSRegistryBusy, XDSRepositoryBusy, XDSRegistryError, XDSRepositoryError, XDSRegistryOutOfResources, XDSRepositoryOutOfResources.

4.18 Receiving Participant Processing Requirements: FHIR Resource Option

Informative: Pushing individual resources across communities is not a typical use case. Usually, rather than simple acceptance and persistence of the resource by the server, such a push will trigger a workflow. This section describes and constrains the acceptable cases.

CONF-276: A FHIR Receiving Participant utilizing the FHIR Resource Option SHALL be constrained to the cases defined in this section, or in another eHealth Exchange profile.

CONF-277: A FHIR Receiving Participant utilizing the FHIR Resource Option SHALL support the following mechanisms for submitting resources:

- create (POST: see https://www.hl7.org/fhir/http.html#create)
- update (PUT: see https://www.hl7.org/fhir/http.html#update)
- batch/transaction (POST /Bundle with POST or PUT entries: see https://www.hl7.org/fhir/http.html#transaction)

Copyright© 2020-2021 All rights reserved.

Informative: Note that the **update as create** (https://www.hl7.org/fhir/http.html#upsert) and **patch** (https://www.hl7.org/fhir/http.html#patch) interactions are not currently supported for pushing content.

4.18.1 Use Case: Patient Correction

Informative: The Patient Correction use case is owned and defined by the HL7 Patient Empowerment Workgroup.

CONF-278: An Initiating Participant utilizing the FHIR Resource Option for the Patient Correction use case SHALL implement the requirements for servers in [HL7 Patient Correction].

4.19 Receiving Participant Processing Requirements: FHIR Messaging Option

This section is reserved.

5 SECURITY CONSIDERATIONS

All messages transacted under this specification SHALL meet or exceed the eHealth Exchange security requirements documented in the Authorization Framework, Messaging Platform, Operational Policies and Procedures, etc. This includes encrypting all messages while at-rest and in-transit and using 2-way-TLS with mutual authentication.

Implementers are encouraged to read the relevant Security Considerations in the IHE ITI TF, specifically [IHE ITI TF-1] 10.7 and [IHE ITI TF-2x] Appendix K.

In the case of cross-author update (an author/organization submits an update to the clinical data from another author/organization—see section 3.14.2), we have described additional provenance requirements to ensure that authorship is traced appropriately. But even before such an update is accepted, the Receiving Participant may want to apply additional checks on the sender, including human review. In this case, the DocumentQueued warning MAY be used to notify the sender of the delay, and the codeContext field MAY be used to explain it in more detail. In general, for cross-author updates to be legitimate, they should be an anticipated part of a use case, for example, multiple authors collaborating on a workflow by successively replacing an IHE XDW document or a FHIR Task with their updates.

As of this writing, FHIR security requirements are in a high degree of flux. The eHealth Exchange will publish a FHIR security requirements specification, as an update to the Authorization Framework specification, that MUST be followed for eHealth Exchange FHIR transactions.

Copyright© 2020-2021 All rights reserved.

6 AUDITING

All messages transacted SHALL meet or exceed the IHE ATNA audit logging requirements found in the respective transactions.

7 OPERATIONAL CONSIDERATIONS

The various IHE and eHx options defined in this specification imply some way for eHx participants to advertise what they support and to discover what others support.

7.1 Directory Settings

CONF-279: If a Receiving Participant supports both the XDR and XCDR Transaction Options, it SHALL represent each as its own Endpoint element in the Directory. These endpoints MAY use the same address URL.

CONF-280: A Receiving Participant that implements the Document Submission specification SHALL use the following settings for the Endpoint in the Directory:

- Endpoint/name/value: “Document Submission”
- Endpoint/extension/extension[url=“Version"]/valueCodeableConcept/coding/value: ”3.0”
- If declaring options:
    - Endpoint/extension/extension[url=“Option"]/valueCodeableConcept/coding/value: ”PatientIdRequired”
    - Endpoint/extension/extension[url=“Option"]/valueCodeableConcept/coding/value: ”PersistsOriginalDocuments”
    - Endpoint/extension/extension[url=“Option"]/valueCodeableConcept/coding/value: ”PersistsClinicalItems”

CONF-281: A Receiving Participant that implements the Document Submission specification with the XDR Transaction Option SHALL use the following settings for the Endpoint in the Directory:

- Endpoint/connectionType/code/value: “ihe-xdr”
- Endpoint/extension/extension[url=“Transaction"]/valueString/value: “XDR ITI-41”
- Endpoint/extension/extension[url=“Actor"]/valueCodeableConcept/coding/value: ”Document Recipient”

Copyright© 2020-2021 All rights reserved.

- If declaring options:
    - Endpoint/extension/extension[url="Option"]/valueCodeableConcept/coding/value: "AcceptsLimitedMetadata"

CONF-282: A Receiving Participant that implements the Document Submission specification with the XCDR Transaction Option SHALL use the following settings for the Endpoint in the Directory:

- Endpoint/connectionType/code/value: "ihe-xcdr"
- Endpoint/extension/extension[url="Transaction"]/valueString/value: "XCDR ITI-80"
- Endpoint/extension/extension[url="Actor"]/valueCodeableConcept/coding/value: "Responding Gateway"

Informative: The connection types are defined in the HL7 value set: http://hl7.org/fhir/valueset-endpoint-connection-type.html. Note that the value "ihe-xcdr" is not yet defined in the HL7 value set referenced by the directory. We will be proposing the additional value.

8 Network-Network Considerations

Some technical gateways as implemented by other networks require that XML-Digital Signatures in the SOAP messages be signed by a PKI private key specific to that network. When such requirements exist, the eHealth Exchange Hub will remove each original XML-Digital Signature and replace it with one compatible with the peer network. Note that this event is audited to preserve the chain of trust from the ultimate sender to the ultimate receiver.

9 eHx Hub Considerations

The Hub may offer the following translation capabilities between the XDR and XCDR flavors. These capabilities will be transparent to Participants.

9.1 Hub Translation: XDR Initiating Participant to XCDR Receiving Participant

In this translation, the following behaviors are performed by the Hub and the Directory:

- For each XCDR Receiving Participant that does not also include an XDR endpoint, eHx staff will create an additional XDR endpoint in the Directory.
- If a request comes in on the added XDR endpoint:

Copyright© 2020-2021 All rights reserved.

    - If it is an ITI-41 request, the Hub will convert it to an ITI-80 request and forward it to the participant's XCDR endpoint.
    - If the ITI-41 request includes a federated Home Community ID, the Hub will return an error.
    - If it is an ITI-80 request, the Hub will return an error.
- For each child of an XCDR Receiving Participant in the Directory that does not have its own native endpoint, eHx staff will create an additional XDR endpoint in the Directory that includes the HCID in the URL.
- If a request comes in on an added child XDR endpoint:
    - If it is an ITI-41 request, the Hub will convert it to an ITI-80 request, extract the HCID from the URL and include it as a federated Home Community ID in the ITI-80 request and forward it to the participant's XCDR endpoint.
    - If the ITI-41 request includes a federated Home Community ID, the Hub will return an error.
    - If it is an ITI-80 request, the Hub will return an error.

9.2 Hub Translation: XCDR Initiating Participant to XDR Receiving Participant

In this translation, the following behaviors are performed by the Hub and the Directory:

- For each XDR Receiving Participant that does not also include an XCDR endpoint, eHx staff will create an additional XCDR endpoint in the Directory.
- If a request comes in on the added XCDR endpoint:
    - If it is an ITI-80 request, the Hub will convert it to an ITI-41 request and forward it to the participant's XDR endpoint.
    - If the ITI-80 request includes a federated Home Community ID, the Hub will look for a child organization in the Directory with that HCID and an XDR endpoint. If it can find one, it will convert the request to an ITI-41 request without a federated Home Community ID and forward it to the child's XDR endpoint. If it cannot find one, the Hub will return an error.
    - If it is an ITI-41 request, the Hub will return an error.

For Hub behaviors that mediate cross-network differences, see the above Network-Network Considerations Section.

Copyright© 2020-2021 All rights reserved.

10 Sequoia Test Tool Considerations

The eHx may leverage the Sequoia Project test platform in the future to provide for semi-automated validation of correct Document Submission transaction sender and/or receiver operations.

11 Roadmap/Backlog

Change	Trigger	Status
Transition period begins	Approval by CC	2021-05-18
New transactions available for trial use by Participants. Participants adopting prior version ensure there are no problems adopting new version.	Approval by CC	2015-05-18
Interim test criteria effective	Approval by CC	Not started
Test implementation Changes deployed as provisional	Approval by CC	Not started
Transition period ends. New specification effective. Prior version 1.0 deprecated.	Approval and date identified by CC	Not started
Test implementation changes transitioned to effective	New specification effective	Not started

Copyright© 2020-2021 All rights reserved.

12 APPENDICES

12.1 Validation Plan

An associated validation plan will confirm the conformance statements in this specification.

12.2 Examples

12.2.1 Document Submission XDR Request Message

```
<?xml version="1.0" encoding="UTF-8"?>
<s:Envelope xmlns:s="http://www.w3.org/2003/05/soap-envelope"
            xmlns:a="http://www.w3.org/2005/08/addressing">
  <s:Header>
    <a:Action s:mustUnderstand="1">
      urn:ihe:iti:2007:ProvideAndRegisterDocumentSet-b
    </a:Action>
    <a:MessageID>urn:uuid:6d296e90-e5dc-43d0-b455-7c1f3eb35d83</a:MessageID>
    <a:ReplyTo>
      <a:Address>http://www.w3.org/2005/08/addressing/anonymous</a:Address>
    </a:ReplyTo>
    <a:To s:mustUnderstand="1">https://reportingagency.org/anEndpoint</a:To>
    <a:From>http://https://generalhospital.org/anEndpoint</a:From>
    <wsse:Security s:mustUnderstand="true">
      <!-- Includes necessary security header information as defined
           in the Messaging Platform Specification -->
    </wsse:Security>
  </s:Header>
```


Copyright© 2020-2021 All rights reserved.

```
<s:Body>
  <xds:ProvideAndRegisterDocumentSetRequest
      xmlns:xds="urn:ihe:iti:xds-b:2007"
      xmlns:lcm="urn:oasis:names:tc:ebxml-regrep:xsd:lcm:3.0"
      xmlns:rim="urn:oasis:names:tc:ebxml-regrep:xsd:rim:3.0"
      xmlns:rs="urn:oasis:names:tc:ebxml-regrep:xsd:rs:3.0">
    <lcm:SubmitObjectsRequest>
      <rim:RegistryObjectList>
        <!-- Note that specifying the entryUUID (ExtrinsicObject/@id)
             allows for submitting corrections later -->
        <rim:ExtrinsicObject id="urn:uuid:c9230bcc-818e-40e5-9df8-
076c5c5d8af9" mimeType="text/xml" objectType="urn:uuid:7edca82f-054d-47f2-
a032-9b2a5b5186c1">
          <rim:Slot name="creationTime">
            <rim:ValueList>
              <rim:Value>20051224</rim:Value>
            </rim:ValueList>
          </rim:Slot>
          <Slot name="hash">
            <ValueList>
              <Value>3278dd4a5b4710bebbc68267a642d12b55394697</Value>
            </ValueList>
          </Slot>
          <Slot name="languageCode">
            <ValueList>
              <Value>en-US</Value>
```


Copyright© 2020-2021 All rights reserved.

```
    </ValueList>
  </Slot>
  <rim:Slot name="languageCode">
    <rim:ValueList>
      <rim:Value>en-us</rim:Value>
    </rim:ValueList>
  </rim:Slot>
  <rim:Slot name="serviceStartTime">
    <rim:ValueList>
      <rim:Value>200412230800</rim:Value>
    </rim:ValueList>
  </rim:Slot>
  <rim:Slot name="serviceStopTime">
    <rim:ValueList>
      <rim:Value>200412230801</rim:Value>
    </rim:ValueList>
  </rim:Slot>
  <Slot name="size">
    <ValueList>
      <Value>381072</Value>
    </ValueList>
  </Slot>
  <rim:Slot name="sourcePatientId">
    <rim:ValueList>
```


Copyright© 2020-2021 All rights reserved.

```
                <rim:Value>ST-
1000^^^&amp;1.3.6.1.4.1.21367.2003.3.9&amp;ISO</rim:Value>

              </rim:ValueList>

            </rim:Slot>

            <rim:Slot name="sourcePatientInfo">

              <rim:ValueList>

                <rim:Value>PID-3|ST-
1000^^^&amp;1.3.6.1.4.1.21367.2003.3.9&amp;ISO</rim:Value>

                <rim:Value>PID-5|Doe^John^^^</rim:Value>

                <rim:Value>PID-7|19560527</rim:Value>

                <rim:Value>PID-8|M</rim:Value>

                <rim:Value>PID-11|100 Main
St^^Metropolis^Il^44130^USA</rim:Value>

              </rim:ValueList>

            </rim:Slot>

            <rim:Name>

              <rim:LocalizedString value="Discharge summary"/>

            </rim:Name>

            <rim:Description/>

            <rim:Classification id="cl01"
classificationScheme="urn:uuid:93606bcf-9494-43ec-9b4e-a7748d1a838d"
classifiedObject="urn:uuid:c9230bcc-818e-40e5-9df8-076c5c5d8af9">

              <rim:Slot name="authorPerson">

                <rim:ValueList>

                  <rim:Value>Gerald Smitty</rim:Value>

                </rim:ValueList>

              </rim:Slot>

              <rim:Slot name="authorInstitution">
```


Copyright© 2020-2021 All rights reserved.

```
                <rim:ValueList>
                  <rim:Value>Cleveland Clinic</rim:Value>
                  <rim:Value>Parma Community</rim:Value>
                </rim:ValueList>
              </rim:Slot>
              <!-- Example of coded value -->
              <rim:Slot name="authorRole">
                <rim:ValueList>
                  <rim:Value>PRF^^^&amp;2.16.840.1.113883.5.90&amp;ISO</rim:V
alue>
                </rim:ValueList>
              </rim:Slot>
              <!-- Example of simple string -->
              <rim:Slot name="authorSpecialty">
                <rim:ValueList>
                  <rim:Value>Cardiology</rim:Value>
                </rim:ValueList>
              </rim:Slot>
            </rim:Classification>
            <rim:Classification id="cl02"
classificationScheme="urn:uuid:41a5887f-8865-4c09-adf7-e362475b143a"
classifiedObject="urn:uuid:c9230bcc-818e-40e5-9df8-076c5c5d8af9"
nodeRepresentation="18842-5">
              <rim:Slot name="codingScheme">
                <rim:ValueList>
                  <rim:Value>2.16.840.1.113883.6.1</rim:Value>
                </rim:ValueList>
```


Copyright© 2020-2021 All rights reserved.

```
            </rim:Slot>

            <rim:Name>

              <rim:LocalizedString value="Discharge summary"/>

            </rim:Name>

          </rim:Classification>

          <rim:Classification id="cl03"
classificationScheme="urn:uuid:f4f85eac-e6cb-4883-b524-f2705394840f"
classifiedObject="urn:uuid:c9230bcc-818e-40e5-9df8-076c5c5d8af9"
nodeRepresentation="N">

            <rim:Slot name="codingScheme">

              <rim:ValueList>

                <rim:Value>2.16.840.1.113883.5.25</rim:Value>

              </rim:ValueList>

            </rim:Slot>

            <rim:Name>

              <rim:LocalizedString value="Normal"/>

            </rim:Name>

          </rim:Classification>

          <rim:Classification id="cl04"
classificationScheme="urn:uuid:a09d5840-386c-46f2-b5ad-9c3699a4309d"
classifiedObject="urn:uuid:c9230bcc-818e-40e5-9df8-076c5c5d8af9"
nodeRepresentation="urn:hl7-org:sdwg:ccda-structuredBody:2.1">

            <rim:Slot name="codingScheme">

              <rim:ValueList>

                <rim:Value>1.3.6.1.4.1.19376.1.2.3</rim:Value>

              </rim:ValueList>

            </rim:Slot>

            <rim:Name>
```


Copyright© 2020-2021 All rights reserved.

```
                <rim:LocalizedString value="C-CDA 2.1 constraints using a
structured body"/>

              </rim:Name>

            </rim:Classification>

            <rim:Classification id="cl05"
classificationScheme="urn:uuid:f33fb8ac-18af-42cc-ae0e-ed0b0bdb91e1"
classifiedObject="urn:uuid:c9230bcc-818e-40e5-9df8-076c5c5d8af9"
nodeRepresentation="73770003">

              <rim:Slot name="codingScheme">

                <rim:ValueList>

                  <rim:Value>2.16.840.1.113883.6.96</rim:Value>

                </rim:ValueList>

              </rim:Slot>

              <rim:Name>

                <rim:LocalizedString value="Emergency department--hospital"/>

              </rim:Name>

            </rim:Classification>

            <rim:Classification id="cl06"
classificationScheme="urn:uuid:cccf5598-8b07-4b77-a05e-ae952c785ead"
classifiedObject="urn:uuid:c9230bcc-818e-40e5-9df8-076c5c5d8af9"
nodeRepresentation="394579002">

              <rim:Slot name="codingScheme">

                <rim:ValueList>

                  <rim:Value>2.16.840.1.113883.6.96</rim:Value>

                </rim:ValueList>

              </rim:Slot>

              <rim:Name>

                <rim:LocalizedString value="Cardiology"/>

              </rim:Name>
```


Copyright© 2020-2021 All rights reserved.

```
            </rim:Classification>

            <rim:Classification id="cl07"
classificationScheme="urn:uuid:f0306f51-975f-434e-a61c-c59651d33983"
classifiedObject="urn:uuid:c9230bcc-818e-40e5-9df8-076c5c5d8af9"
nodeRepresentation="59258-4">

              <rim:Slot name="codingScheme">

                <rim:ValueList>

                  <rim:Value>2.16.840.1.113883.6.1</rim:Value>

                </rim:ValueList>

              </rim:Slot>

              <rim:Name>

                <rim:LocalizedString value="Emergency department Discharge
summary"/>

              </rim:Name>

            </rim:Classification>

            <rim:ExternalIdentifier id="ei01"
registryObject="urn:uuid:c9230bcc-818e-40e5-9df8-076c5c5d8af9"
identificationScheme="urn:uuid:58a6f841-87b3-4a3e-92fd-a8ffeff98427"
value="SELF-5^^^&amp;1.3.6.1.4.1.21367.2005.3.7&amp;ISO">

              <rim:Name>

                <rim:LocalizedString value="XDSDocumentEntry.patientId"/>

              </rim:Name>

            </rim:ExternalIdentifier>

            <rim:ExternalIdentifier id="ei02"
registryObject="urn:uuid:c9230bcc-818e-40e5-9df8-076c5c5d8af9"
identificationScheme="urn:uuid:2e82c1f6-a085-4c72-9da3-8640a32e42ab"
value="1.3.6.1.4.1.21367.2005.3.9999.32">

              <rim:Name>

                <rim:LocalizedString value="XDSDocumentEntry.uniqueId"/>

              </rim:Name>
```


Copyright© 2020-2021 All rights reserved.

```
            </rim:ExternalIdentifier>

          </rim:ExtrinsicObject>

          <rim:RegistryPackage id="SubmissionSet01">

            <rim:Slot name="submissionTime">

              <rim:ValueList>

                <rim:Value>20041225235050</rim:Value>

              </rim:ValueList>

            </rim:Slot>

            <rim:Name>

              <rim:LocalizedString value="Hospital Stay"/>

            </rim:Name>

            <rim:Description/>

            <rim:Classification id="cl08"
classificationScheme="urn:uuid:a7058bb9-b4e4-4307-ba5b-e3f0ab85e12d"
classifiedObject="SubmissionSet01">

              <rim:Slot name="authorPerson">

                <rim:ValueList>

                  <rim:Value>Sherry Dopplemeyer</rim:Value>

                </rim:ValueList>

              </rim:Slot>

              <rim:Slot name="authorInstitution">

                <rim:ValueList>

                  <rim:Value>Cleveland Clinic</rim:Value>

                  <rim:Value>Berea Community</rim:Value>

                </rim:ValueList>

              </rim:Slot>
```


Copyright© 2020-2021 All rights reserved.

```
                <rim:Slot name="authorRole">
                  <rim:ValueList>
                    <rim:Value>Primary Surgeon</rim:Value>
                  </rim:ValueList>
                </rim:Slot>
                <rim:Slot name="authorSpecialty">
                  <rim:ValueList>
                    <rim:Value>Orthopedic</rim:Value>
                  </rim:ValueList>
                </rim:Slot>
              </rim:Classification>
              <rim:Classification id="cl09"
classificationScheme="urn:uuid:aa543740-bdda-424e-8c96-df4873be8500"
classifiedObject="SubmissionSet01" nodeRepresentation="EMER">
                <rim:Slot name="codingScheme">
                  <rim:ValueList>
                    <rim:Value>2.16.840.1.113883.5.4</rim:Value>
                  </rim:ValueList>
                </rim:Slot>
                <rim:Name>
                  <rim:LocalizedString value="Emergency"/>
                </rim:Name>
              </rim:Classification>
              <rim:ExternalIdentifier id="ei03"
registryObject="SubmissionSet01" identificationScheme="urn:uuid:96fdda7c-
d067-4183-912e-bf5ee74998a8" value="1.3.6.1.4.1.21367.2005.3.9999.33">
                <rim:Name>
```


Copyright© 2020-2021 All rights reserved.

```
                <rim:LocalizedString value="XDSSubmissionSet.uniqueId"/>

              </rim:Name>

            </rim:ExternalIdentifier>

            <rim:ExternalIdentifier id="ei04"
registryObject="SubmissionSet01" identificationScheme="urn:uuid:554ac39e-
e3fe-47fe-b233-965d2a147832" value="3670984664">

              <rim:Name>

                <rim:LocalizedString value="XDSSubmissionSet.sourceId"/>

              </rim:Name>

            </rim:ExternalIdentifier>

            <rim:ExternalIdentifier id="ei05"
registryObject="SubmissionSet01" identificationScheme="urn:uuid:6b5aea1a-
874d-4603-a4bc-96a0a7b38446" value="SELF-
5^^^&amp;1.3.6.1.4.1.21367.2005.3.7&amp;ISO">

              <rim:Name>

                <rim:LocalizedString value="XDSSubmissionSet.patientId"/>

              </rim:Name>

            </rim:ExternalIdentifier>

          </rim:RegistryPackage>

          <rim:Classification id="cl10" classifiedObject="SubmissionSet01"
classificationNode="urn:uuid:a54d6aa5-d40d-43f9-88c5-b4633d873bdd"/>

          <rim:Association id="as01" associationType="HasMember"
sourceObject="SubmissionSet01" targetObject="urn:uuid:c9230bcc-818e-40e5-
9df8-076c5c5d8af9">

            <rim:Slot name="SubmissionSetStatus">

              <rim:ValueList>

                <rim:Value>Original</rim:Value>

              </rim:ValueList>

            </rim:Slot>
```


Copyright© 2020-2021 All rights reserved.

```
          </rim:Association>

        </rim:RegistryObjectList>

      </lcm:SubmitObjectsRequest>

      <xds:Document id="urn:uuid:c9230bcc-818e-40e5-9df8-
076c5c5d8af9">UjBsR09EbGhjZ0dTQUxNQUFBUUNBRU1tQ1p0dU1GUXhEUzhi</xds:Document>

    </xds:ProvideAndRegisterDocumentSetRequest>

  </s:Body>

</s:Envelope>
```

12.2.2 Document Submission XDR Response Message

```
<?xml version="1.0" encoding="UTF-8"?>

<s:Envelope xmlns:s="http://www.w3.org/2003/05/soap-envelope"

            xmlns:a="http://www.w3.org/2005/08/addressing">

  <s:Header>

    <a:Action s:mustUnderstand="1">

      urn:ihe:iti:2007:ProvideAndRegisterDocumentSet-bResponse

    </a:Action>

    <a:RelatesTo>urn:uuid:6d296e90-e5dc-43d0-b455-7c1f3eb35d83</a:RelatesTo>

  </s:Header>

  <s:Body>

    <rs:RegistryResponse

      status="urn:oasis:names:tc:ebxml-regrep:ResponseStatusType:Success"

      xmlns:rs="urn:oasis:names:tc:ebxml-regrep:xsd:rs:3.0"/>

  </s:Body>

</s:Envelope>
```


Copyright© 2020-2021 All rights reserved.

12.2.3 Document Submission XCDR Request Message

```
<?xml version="1.0" encoding="UTF-8"?>
<s:Envelope xmlns:s="http://www.w3.org/2003/05/soap-envelope"
            xmlns:a="http://www.w3.org/2005/08/addressing"
            xmlns:xdr="urn:ihe:iti:xdr:2014">
  <s:Header>
    <a:Action s:mustUnderstand="1">
      urn:ihe:iti:2015:CrossGatewayDocumentProvide
    </a:Action>
    <xdr:homeCommunityBlock>
      <xdr:homeCommunityId>urn:oid:1.2.3.4.5.6.2333.23</xdr:homeCommunityId>
    </xdr:homeCommunityBlock>
    <a:MessageID>urn:uuid:6d296e90-e5dc-43d0-b455-7c1f3eb35d83</a:MessageID>
    <a:ReplyTo>
      <a:Address>http://www.w3.org/2005/08/addressing/anonymous</a:Address>
    </a:ReplyTo>
    <a:To s:mustUnderstand="1">https://reportingagency.org/anEndpoint</a:To>
    <a:From>http://https://generalhospital.org/anEndpoint</a:From>
    <wsse:Security s:mustUnderstand="true">
      <!-- Includes necessary security header information as defined
           in the Messaging Platform Specification -->
    </wsse:Security>
  </s:Header>
  <s:Body>
    <xds:ProvideAndRegisterDocumentSetRequest
```


Copyright© 2020-2021 All rights reserved.

```
        xmlns:xds="urn:ihe:iti:xds-b:2007"
        xmlns:lcm="urn:oasis:names:tc:ebxml-regrep:xsd:lcm:3.0"
        xmlns:rim="urn:oasis:names:tc:ebxml-regrep:xsd:rim:3.0"
        xmlns:rs="urn:oasis:names:tc:ebxml-regrep:xsd:rs:3.0">
      <lcm:SubmitObjectsRequest>
        <rs:RequestSlotList>
          <rim:Slot name="homeCommunityId">
            <rim:ValueList>
              <rim:Value>urn:oid:1.2.3.4.5.6.2333.23</rim:Value>
            </rim:ValueList>
          </rim:Slot>
        </rs:RequestSlotList>
        <rim:RegistryObjectList>
          <!-- Registry Metadata goes here -->
        </rim:RegistryObjectList>
      </lcm:SubmitObjectsRequest>
      <xds:Document id="urn:uuid:c9230bcc-818e-40e5-9df8-076c5c5d8af9">
        <!-- Document binary goes here -->
      </xds:Document>
    </xds:ProvideAndRegisterDocumentSetRequest>
  </s:Body>
</s:Envelope>
```

12.2.4 Document Submission XCDR Response Message

```
<?xml version="1.0" encoding="UTF-8"?>
<s:Envelope xmlns:s="http://www.w3.org/2003/05/soap-envelope"
```


Copyright© 2020-2021 All rights reserved.

```
    xmlns:a="http://www.w3.org/2005/08/addressing">
  <s:Header>
    <a:Action s:mustUnderstand="1">
      urn:ihe:iti:2015:CrossGatewayDocumentProvideResponse
    </a:Action>
    <a:RelatesTo>urn:uuid:6d296e90-e5dc-43d0-b455-7c1f3eb35d83</a:RelatesTo>
  </s:Header>
  <s:Body>
    <rs:RegistryResponse
      status="urn:oasis:names:tc:ebxml-regrep:ResponseStatusType:Success"
      xmlns:rs="urn:oasis:names:tc:ebxml-regrep:xsd:rs:3.0"
  </s:Body>
</s:Envelope>
```

12.3 Open issues

1.

12.4 Closed issues

1. Which flavors of Push should we adopt: XDR, XCDR, XDM, XDS.b?
    a. Resolution: Because there was a need to route to federated systems, XCDR was adopted. To maintain backwards compatibility with the prior version of the specification, XDR was adopted. There were no use cases that required XDM or XDS.b, however, both were analyzed, and some of the consistency checking of XDS.b was added.
2. There is no Metadata-Limited Document Source actor or option for the XCDR Initiating Gateway. Is this something we need to add?
    a. Resolution: No use case for this.
3. Is there a need for asynchronous push? The prior version had a deferred mechanism, but it wasn't based on an IHE mechanism. Only one option in the current IHE ITI TF is available: AS4.

Copyright© 2020-2021 All rights reserved.

    a. Resolution: No use case for this.
4. Should we define any persistence requirements?
    a. Resolution: Yes, added conditional requirements depending on whether and how a system persisted submissions.
5. Should we define any provenance requirements?
    a. Resolution: Yes, examined provenance in the context of push, and defined specific cases and provenance info that needs to be retained. Leaned heavily on work done by HL7 and IHE.

Copyright© 2020-2021 All rights reserved.

12.5 Push Comparison Table

The following table was created to assist in the writing of this specification. It is a cross-reference of detailed functionality across the various flavors of push:

- Columns B through E reflect the IHE ITI push profiles in the 2019 Technical Framework, which are being referenced by this specification
- Columns F and G reflect the eHx 2011 Document Submission specification and its underlying specification, the IHE ITI XDR profile in the 2009 Technical Framework
- Column H reflects the decisions made for this specification. Contents will be either specific notes or the following:
    - <conformance word> IHE: this spec applies the IHE requirement implicitly
    - <conformance word> eHx: this spec adds an explicit requirement
    - MAY-UND: this spec does not utilize this functionality, so any behavior is undefined
- Columns I and J trace the functionality to specific sections
- Column K analyzes compatibility of the 2019 IHE XCDR profile with the TEFCA QTF
- Cells are shaded pink to indicate greater relative importance

Below is the table as an embedded Excel file followed by the table in Word format. To view the embedded table in Excel, assuming all software prerequisites are met, double click on the below image. (This only works for the Word format version of this file. For the PDF version please see the next few pages.)

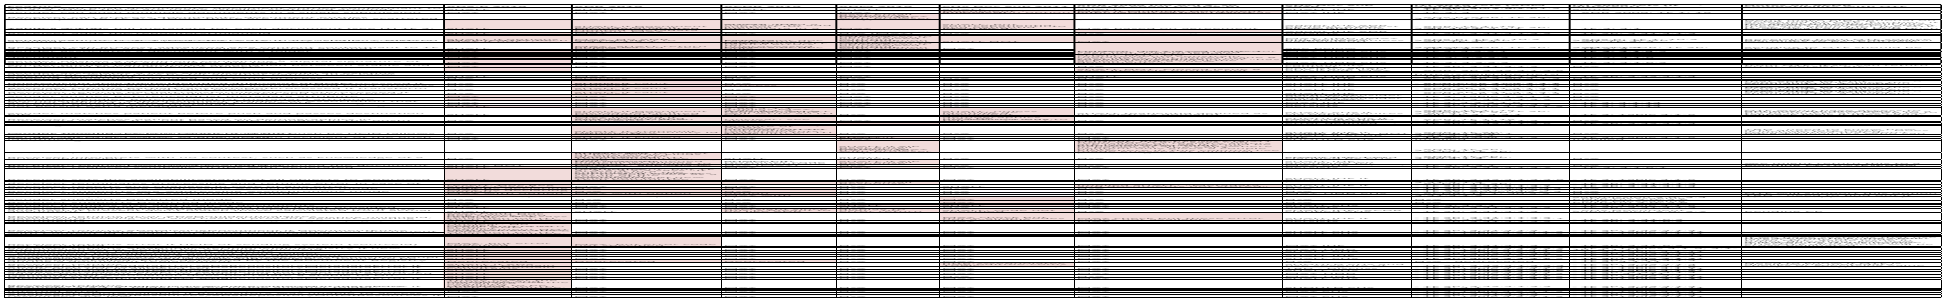

Copyright© 2020-2021 All rights reserved.

Feature	XDS.b 2019	XDR 2019	XCDR 2019	XDM 2019	eHx Doc Sub 2011	XDR 2009 (no Vol 2, so mostly equal to XDS.b)	eHx Doc Sub 2020	ITI 2019 refs (XDM omitted)	eHx Doc Sub 2011, ITI 2009 refs (if different)	Notes on XCDR compatibility with QTF
Basic metadata: documents, submission set										
Sender can push documents, document entries conformant to TF 3: 4 (Note there are many "shalls" scattered around)	SHALL	SHALL	SHALL	SHALL	SHALL	SHALL	SHALL IHE	- TF 2b: 3.41.4.1.2 - XCDR: 2b 3.80.4.1.2 - TF-3: 4	Same	
Sender can push multiple documents in a single submission	MAY	MAY	MAY	MAY	MAY (Multiple Document Submission option included)	MAY if Multiple Document Submission option (this was later removed from XDR)	MAY IHE	- TF-3: 4	- XDR 2009: TF 1: 15	
Receiver can process documents, document entries conformant to TF 3: 4 (Note there are many "shalls" scattered around)	SHALL	SHALL	SHALL	SHALL if ZIP over Email Response Option; MAY otherwise (i.e. render with browser only)	SHALL	SHALL	SHALL IHE	- XDS.b/XDR: TF 2b: 3.41.4.1.3 - TF-3: 4	Same	
Sender pushes full metadata for each document (see ITI TF 3: Table 4.3.1-3)	SHALL	SHALL if Document Source; MAY if Metadata-Limited Document Source	SHALL. IHE would need a use case to add support for limited metadata in a CP.	MAY	SHALL, but sourcePatientId, sourcePatientInfo and patientId MAY be de-identified	MAY based on agreement	Adopt IHE XDR options, suggest CP to XCDR to allow omitting patientId	- XDR: TF 1: 15 - XDS.b/XDR: TF 2b: 3.41.4.1.2 - All: TF-3: 4.3.1	- TF-3: 4.1	Don't make PID optional on XCDR until ITI CP is accepted by QTF. Further, if we want full compatibility for all pushes, we will have to disallow limited metadata pushes for XDR if they are bound for the QTF. If we do, then the Hub could still convert XDR push to XCDR to go out to QTF.
Receiver handles partial metadata	SHALL NOT	SHALL if Accepts Limited Metadata Option; MAY otherwise	MAY	MAY	Unclear	MAY based on agreement	Adopt IHE XDR options	- TF 2b: 3.41.4.1.3	N/A	Ok for receivers
Receiver adds the repositoryUniqueId	SHALL	MAY	MAY	N/A (MAY)	N/A	N/A	SHALL IHE if persists	- TF 2b: 3.41.4.1.3.2	- TF 2b: 3.41.4.1.3	
Advanced metadata: associations, folders										
Sender can push advanced XDS content: associations, folders conformant to TF 3: 4 (Note there are many "shalls" scattered around)	SHALL if options declared for each advanced type; MAY if not	MAY. No XDR counterpart to options in 3.41.4.1.2.1.	MAY for Folders or doc replacement	MAY, but references to existing objects can't be resolved without external coordination with receiver	SHALL NOT	MAY	MAY-UND (MAY but behavior undefined unless higher level agreement)	- XDS.b: TF 1: 10.2 - XDS.b: TF 2b: 3.41.4.1.2.1 - All: TF-3: 4	- XDS.b: TF 1: 10.2 - XDS.b: TF 2b: 3.41.4.1.2.1 - All: TF-3: 4	Receiving from QTF source passing these is probably ok, as we require all Partial warnings to be returned.
Receiver can process advanced XDS content conformant to TF 3: 4 (Note there are many "shalls" scattered around)	SHALL	MAY; SHALL return errors for unprocessed - see below	MAY but unclear; SHALL return errors for unprocessed Folders or doc replacement only	MAY, but references to existing objects can't be resolved without external coordination with receiver	MAY		See below	- XDS.b/XDR: TF 2b: 3.41.4.1.3 - TF-3: 4	- XDS.b/XDR: TF 2b: 3.41.4.1.4 - TF-3: 4	Sending to QTF should be ok as well.
Sender pushes a folder	MAY	MAY	MAY	MAY			MAY-UND eHx	- TF 3: 4.2.2.1	- TF-3: 4.1.4.2 - TF 3: 4.1.5	
Sender associates a document with a folder	MAY	MAY	MAY	MAY		Unclear. 15.2.3 says MAY for Folders, but closed issue 9 says not supported and	MAY-UND eHx	- TF 3: 4.2.2.1	- TF-3: 4.1.4.2 - TF 3: 4.1.5	
Receiver persists folders and associations	SHALL	MAY	MAY	MAY			MAY-UND eHx	- TF 2b: 3.42.4.1.3	- TF 2b: 3.42.4.1.4	

Feature	XDS.b 2019	XDR 2019	XCDR 2019	XDM 2019	eHx Doc Sub 2011	XDR 2009 (no Vol 2, so mostly equal to XDS.b)	eHx Doc Sub 2020	ITI 2019 refs (XDM omitted)	eHx Doc Sub 2011, ITI 2009 refs (if different)	Notes on XCDR compatibility with QTF
Receiver updates Folder.lastUpdateTime if change to folder	SHALL	MAY	MAY	MAY		recipient has to ignore.	MAY-UND IHE	- TF 2b: 3.42.4.1.3.6	- TF 3: Table 4.1-7	
Sender pushes addendum to a document	MAY	MAY	MAY	MAY		SHALL NOT. No requirements constrain metadata, but closed issue 8 says lifecycle pushes not supported and recipient "has to send a negative acknowledgment if the action is not a “new document”."	SHALL eHx	- TF-3: 4.2.2.2	- TF-3: 4.1.6	
Sender pushes replacement of a document	MAY	MAY	MAY	MAY			SHALL eHx	- TF-3: 4.2.2.2	- TF-3: 4.1.6	
Sender pushes transformation of a document	MAY	MAY	MAY	MAY			MAY-UND eHx	- TF-3: 4.2.2.2	- TF-3: 4.1.6	
Sender pushes transformation and replacement of a document	MAY	MAY	MAY	MAY			MAY-UND eHx	- TF-3: 4.2.2.2	- TF-3: 4.1.6	
Receiver persists document relationship semantics	SHALL	MAY	MAY	MAY			SHALL IHE if persists	- TF 2b: 3.42.4.1.3	- TF 2b: 3.42.4.1.4	
Sender pushes a document containing a digital signature of another document and relates the two	MAY	MAY	MAY	MAY			MAY-UND eHx	- TF-3: 4.2.2.2	- TF-3: 4.1.6.2	
Sender pushes IsSnapshotOf association between Stable and On-Demand DocumentEntry objects	MAY	MAY	MAY	MAY		N/A	SHALL NOT eHx - this makes no sense in a cross-community context	- TF 2b: 3.41.4.1.2	N/A	Unlikely to get one of these from QTF. If so, can return error, persist, etc.
Receiver persists On-Demand snapshot semantics	SHALL	MAY	MAY	MAY		N/A	SHALL NOT eHx	- TF 2b: 3.42.4.1.3	N/A	
Receiver ensures document replacement semantics	SHALL	MAY	MAY	MAY		SHALL NOT. Closed issue 8 says lifecycle pushes not supported.	SHALL IHE if persists; SHALL eHx if any info persisted; MAY otherwise	- XDR: 2b 3.42.4.1.3.5 - XCDR: 2b 3.80.4.1.3	- TF 3: 4.1.6.1	Compatible
Receiver persists SubmissionSets and associations	SHALL	MAY	MAY	MAY		SHALL	MAY-UND eHx	- TF 2b: 3.42.4.1.3	- TF 2b: 3.42.4.1.4	
Sender associates a new SubmissionSet with an existing document, which may be for another patient (use case: mother and child birth records)	MAY	MAY	MAY	MAY		MAY	MAY-UND eHx	- TF 3: 4.2.2.1.1 - https://wiki.ihe.net/index.php/XDS-FAQ_2	- TF 3: 4.1.4.2	
Receiver handles and persists a new SubmissionSet with an existing document	SHALL	MAY	MAY	MAY		SHALL	MAY-UND eHx	- TF 2b: 3.42.4.1.3	- TF 2b: 3.42.4.1.4	
Receiver returns PartialAppendContentNotProcessed if append semantics included and can't process	N/A	SHALL if can't process	N/A	N/A	N/A	N/A	SHALL IHE	- XDR 2b 3.41.4.1.3.1 - XCDR: 2b 3.80.4.1.3	N/A	
Receiver returns PartialFolderContentNotProcessed if folder semantics included and can't process	N/A	SHALL if can't process	SHALL if can't process	N/A	N/A	N/A	SHALL IHE	- XDR 2b 3.41.4.1.3.1 - XCDR: 2b 3.80.4.1.3	N/A	
Receiver returns PartialRelationshipContentNotProcessed if relationship semantics included and can't process	N/A	SHALL if can't process	N/A	N/A	N/A	N/A	SHALL IHE	- XDR 2b 3.41.4.1.3.1 - XCDR: 2b 3.80.4.1.3	N/A	We might return warning that XCDR IG isn't expecting, but since the error code is defined, shouldn't be a problem.
Receiver returns PartialReplaceContentNotProcessed if document replacement semantics included and can't process	N/A	SHALL if can't process	SHALL if can't process	N/A	N/A	N/A	SHALL IHE	- XDR 2b 3.41.4.1.3.1 - XCDR: 2b 3.80.4.1.3	N/A	
Receiver returns PartialTransformNotProcessed if transform semantics included and can't process	N/A	SHALL if can't process	N/A	N/A	N/A	N/A	SHALL IHE	- XDR 2b 3.41.4.1.3.1 - XCDR: 2b 3.80.4.1.3	N/A	We might return warning that XCDR IG isn't expecting, but since the error code is defined, shouldn't be a problem.
Receiver returns PartialTransformReplaceNotProcessed if transform&replace semantics included and can't process	N/A	SHALL if can't process	N/A	N/A	N/A	N/A	SHALL IHE	- XDR 2b 3.41.4.1.3.1 - XCDR: 2b 3.80.4.1.3	N/A	We might return warning that XCDR IG isn't expecting, but since the error code is defined, shouldn't be a problem.
Sender pushes ReferenceIdList metadata attribute	MAY	MAY	MAY	MAY	N/A	N/A	MAY-UND unless want to consider encounter persistence	- TF 3: Table 4.3.1-3	N/A	

Feature	XDS.b 2019	XDR 2019	XCDR 2019	XDM 2019	eHx Doc Sub 2011	XDR 2009 (no Vol 2, so mostly equal to XDS.b)	eHx Doc Sub 2020	ITI 2019 refs (XDM omitted)	eHx Doc Sub 2011, ITI 2009 refs (if different)	Notes on XCDR compatibility with QTF
Receiver handles ReferenceIdList metadata attribute	MAY	MAY	MAY	MAY	N/A	N/A	MAY-UND unless want to consider encounter persistence	- TF 2b: 3.42.4.1.3.2	N/A	
Sender pushes extra metadata: ebRIM Slots on any DocumentEntry, SubmissionSet, Folder, or Association that are not defined in the Technical Framework	MAY	MAY	MAY	MAY	N/A	N/A	MAY IHE	- TF 3: 4.2.3.1.6	- TF 3: 4.1.14	
Receiver handles extra metadata without returning error	SHALL	SHALL	SHALL	MAY	N/A	N/A	SHALL IHE	- TF 3: 4.2.3.1.6	- TF 3: 4.1.14	
Receiver persists extra metadata	MAY	MAY	MAY	MAY	N/A	N/A	MAY IHE if persists	- TF 2b: 3.42.4.1.3.2	- TF 3: 4.1.14	
Receiver returns XDSExtraMetadataNotSaved warning if does not persist extra metadata	SHALL	MAY	MAY	N/A	N/A	N/A	SHALL IHE if persists	- TF 2b: 3.42.4.1.3.2 - TF 3: Table 4.2.4.1-2	- TF 3: 4.1.14	
Patient: subject of the push										
Sender matches patient before push and passes destination PID	SHALL	SHALL if Document Source; MAY if Metadata-Limited Document Source	Unclear: 3.80.4.1.1 implies MAY, 3.80.4.1.2 specifies SHALL through Vol 3 requirements. Will write an ITI CP.	N/A	SHALL, unless pushing de-identified document	SHALL (patientId defined as in registry's domain)	SHALL eHx if receiver requires via options	- XDR: TF 1: 15 - XDS.b/XDR: TF 2b: 3.41.4.1.2 - XCDR: TF 2b: 3.80.4.1 - All: TF-3: 4.3.1	- TF 3: Table 4.1-5	QTF won't have access to our custom directory options, but shouldn't be a problem. Expect XCDR IG to try to match, and we have fallback logic if they can't.
Receiver returns XDSUnknownPatientId if PID included and doesn't match a known patient	SHALL	MAY	MAY	N/A	SHALL	SHALL	MAY eHx	- TF 2b: 3.42.4.1.3.3.2	- TF 2b: 3.42.4.1.4 - TF 3: Table 4.1-11	
Receiver matches patient based on demographics in metadata or CDA (i.e. not by the PID included in the push)	N/A	MAY implied if Accepts Limited Metadata Option and no PID; N/A otherwise	Unclear: see above	MAY	Unclear: PID and demographics may be pseudonymous; no defined way to tell	N/A	SHOULD eHx if can't match PID (we make XDR match XCDR)	- XDR: TF 1: 15 - TF 2b: 3.41.4.1.3	N/A	
Sender pushes content for more than one patient	SHALL NOT	SHALL NOT	SHALL NOT	MAY	SHALL NOT	SHALL NOT	SHALL NOT IHE	- TF 2b: 3.41.4.1.2	- TF 2b: 3.41.1	
Receiver can process content for more than one patient	SHALL NOT	N/A (MAY)	SHALL NOT	MAY	SHALL NOT	SHALL NOT	SHALL NOT eHx; error defined below	- TF 2b: 3.42.4.1.3.3.2 - TF 2b: 3.42.4.1.3.5 - TF 3: Table 4.2.4.1-2	- TF 2b: 3.42.4.1.4 - TF 3: Table 4.1-11	
Routing										
Sender specifies system recipient of the push by HCID (allows routing beyond the immediate recipient)	N/A	SHALL if Transmit Home Community Id Option, N/A otherwise	This is the key feature of XCDR. However, unclear: SHALL implied but no normative requirements. Question for ITI Tech Committee.	N/A	N/A	N/A	SHALL IHE if XCDR, SHALL NOT if XDR eHx	- XDR: TF 2b: 3.41.4.1.2.2 - XCDR: 3.80.1	N/A	TBD - need to know how QTF directory works. Does it include federated communities that don't have endpoints? Who can we ask?
Receiver routes to system recipient if identified	N/A	MAY	MAY	N/A	N/A	N/A	SHALL if XCDR, N/A if XDR eHx	- XDR: TF 2b: 3.41.4.1.3.1 - XCDR: 3.80.1	N/A	
Sender specifies human or organization recipient of the push by: person or org name, phone, email	MAY	SHALL if known	MAY	SHALL if known	MAY	MAY	IHE	- TF 2b: 3.41.4.1.2 - TF-3: 4.3.1	- TF 3: Table 4.1-5	
Receiver routes to human or organization recipient if identified	MAY	MAY	MAY	MAY	MAY	MAY	SHOULD eHx	Implied	Implied	
Processing										

Feature	XDS.b 2019	XDR 2019	XCDR 2019	XDM 2019	eHx Doc Sub 2011	XDR 2009 (no Vol 2, so mostly equal to XDS.b)	eHx Doc Sub 2020	ITI 2019 refs (XDM omitted)	eHx Doc Sub 2011, ITI 2009 refs (if different)	Notes on XCDR compatibility with QTF
Receiver processes submission	SHALL	SHALL	SHALL	SHALL if ZIP over Email Response Option; MAY otherwise (i.e. render with browser only)	SHALL	Unclear. XDR has no Vol 2 content; it just reuses ITI-41 from XDS.b and says "no repository or registry actors are involved". So it could be interpreted to imply equivalent behavior, no behavior at all, or anywhere in between. The Vol 1 content only gives a couple hints. Elsewhere in this column I've made guesses at minimal processing.	SHALL IHE	- XDR: TF 2b: 3.41.4.1.3 - XCDR: TF 2b: 3.80.4.1.3	- TF 1: 15	
Receiver interprets with no context, such as knowledge of a prior submission	N/A	SHALL, but impossible to meet unless some kinds of pushes are disallowed or constrained. Question for ITI Tech Committee.	SHALL	SHALL implied	N/A	N/A	SHALL IHE, but constrained and clarified by eHx	- XDR: TF 2b: 3.41.4.1.3.1 - XCDR: TF 2b: 3.80.4.1.3	N/A	
Receiver returns response only after full processing	SHALL	SHALL, but response code DocumentQueued (for manual matching) allows patient matching to be deferred.	SHALL but unclear if response code DocumentQueued is allowed	SHALL if ZIP over Email Response Option	SHALL	SHALL	SHALL IHE, except when queueing for later match (eHx)	- TF 2b: 3.41.4.2	Same	We might return this to a QTF IG. Should be ok.
Receiver returns DocumentQueued if queued the document for future manual matching to a patient	SHALL NOT	Unclear, appears to be MAY. See above.	N/A (MAY)	N/A	N/A	N/A	MAY eHx	- TF-3: 4.3.1	N/A	Pending CP to ask if XDR warning can be returned
Receiver replaces symbolic UUIDs with generated	SHALL	MAY	MAY	MAY	MAY	MAY	SHALL IHE if persists	- TF 2b: 3.42.4.1.3.7	- TF-3: 4.1.6.1	
Receiver sets the availabilityStatus of all objects to Approved	SHALL	Unclear. No normative requirement, but TF 3: 4.2.3.2.2 says "If present in a submission, the submitted value is ignored. It is always set to Approved as a result of the successful submission of new documents." Question for ITI Tech Committee.	MAY	N/A	MAY	MAY	SHALL IHE if persists	- TF 2b: 3.42.4.1.3.3.5	- TF 3: Table 4.1-5	
Receiver persists documents, document entries for future query	SHALL	MAY	MAY	N/A	MAY	MAY	SHALL eHx if persists	- TF 2b: 3.41.4.1.3.2 - TF 2b: 3.42.4.1.3.1	- TF 2b: 3.41.4.1.3 - TF 2b: 3.42.4.1.4	
Receiver terminates processing and returns error in case of error	SHALL	SHALL	SHALL	SHALL if ZIP over Email Response Option	MAY	MAY	SHALL IHE	- XDR: TF 2b: 3.41.4.1.3 - XCDR: TF 2b: 3.80.4.1.3	- TF 2b: 3.41.4.1.3 - TF 2b: 3.42.4.1.4	

Feature	XDS.b 2019	XDR 2019	XCDR 2019	XDM 2019	eHx Doc Sub 2011	XDR 2009 (no Vol 2, so mostly equal to XDS.b)	eHx Doc Sub 2020	ITI 2019 refs (XDM omitted)	eHx Doc Sub 2011, ITI 2009 refs (if different)	Notes on XCDR compatibility with QTF
Receiver reverts any changes in case of any error	Doc Repo MAY revert, Doc Reg SHALL revert	MAY	MAY	N/A	SHALL	Implied SHALL: "Metadata and documents not stored"	SHALL IHE if persists	- TF 2b: 3.41.4.1.3.2 - TF 2b: 3.42.4.1.3.1	- TF 2b: 3.41.4.1.3 - TF 2b: 3.42.4.1.4	
Sender supports WS-Addressing Asynchronous mode	MAY by declaring option	N/A	N/A	N/A	N/A	N/A	N/A	- TF 1: 10.2 - TF 2b: 3.41.4.1.2.1 - TF-2x: Appendix V.3	N/A	
Receiver supports WS-Addressing Asynchronous mode	MAY by declaring option	N/A	N/A	N/A	N/A	N/A	N/A	- TF 1: 10.2 - TF 2b: 3.41.4.1.2.1 - TF-2x: Appendix V.3	N/A	
Sender supports AS4 Asynchronous mode	MAY with option	MAY with option	MAY with option	N/A	N/A	N/A	MAY-UND eHx	- XDS.b/XDR: TF 2b: 3.41.4.1.2 - XCDR: TF 1: 40.2.2	N/A	TBD - need to find out if QTF will support this option
Receiver supports AS4 Asynchronous mode	MAY with option	MAY with option	MAY with option	N/A	N/A	N/A	MAY-UND eHx	- XDS.b/XDR: TF 2b: 3.41.4.1.2 - XCDR: TF 1: 40.2.2	N/A	TBD - need to find out if QTF will support this option
Sender supports Deferred mode	N/A	N/A	N/A	N/A	MAY	N/A	N/A	N/A	- eHx Document Submission 2.0: 3.4 - eHx Messaging Platform 3.0: 3.6	
Receiver supports Deferred mode	N/A	N/A	N/A	N/A	MAY	N/A	N/A	N/A	- eHx Document Submission 2.0: 3.4 - eHx Messaging Platform 3.0: 3.6	
Validating: size and hash										
Sender includes size and hash of documents	MAY	MAY	MAY	SHALL	SHALL	MAY	SHALL eHx	- TF 2b: 3.41.4.1.2 - TF-3: 4.3.1	- eHx Document Submission 2.0: 3.5 - TF 3: Table 4.1-5	
Receiver adds and persists size and hash if not provided	SHALL	MAY	MAY	N/A	N/A	MAY	SHALL IHE if persists	- TF 2b: 3.41.4.1.3.2	- TF 2b: 3.41.4.1.3	
Receiver returns error XDSRepositoryMetadataError if size and hash if present don't match those calculated from actual document	SHALL	SHALL	MAY, but this may be oversight since this is meant to be like XDR	SHALL detect error and display to user - no error return	SHALL; code not specified	N/A	SHALL IHE; XCDR clarified eHx	- TF 2b: 3.41.4.1.3	- eHx Document Submission 2.0: 2.7	Pending CP
Receiver returns error XDSNonIdenticalSize/ XDSNonIdenticalHash if new document for existing uniqueId, based on checking size and hash if present against existing documents	SHALL, but unclear. Vol 2b says both Doc Repo and Doc Reg check and return, Vol 3 says only Doc Reg returns. Question for ITI Tech Committee.	MAY	MAY	N/A	MAY check but these error codes not available	MAY check but these error codes not available	SHALL IHE if persists	- TF 2b: 3.41.4.1.3.2 - TF 2b: 3.42.4.1.3.3.1 - TF-3: 4.3.1	- TF 2b: 3.41.6.2	
Receiver allows same document for existing uniqueId, based on checking size and hash if present against existing documents	SHALL	MAY	MAY	N/A	MAY	MAY	SHALL IHE if persists	- TF 2b: 3.41.4.1.3.2	- TF 2b: 3.41.6.2	
Validating: documents vs metadata										

Feature	XDS.b 2019	XDR 2019	XCDR 2019	XDM 2019	eHx Doc Sub 2011	XDR 2009 (no Vol 2, so mostly equal to XDS.b)	eHx Doc Sub 2020	ITI 2019 refs (XDM omitted)	eHx Doc Sub 2011, ITI 2009 refs (if different)	Notes on XCDR compatibility with QTF
Receiver returns XDSMissingDocument if DocumentEntry exists in metadata with no corresponding attached document	Unclear. No explicit requirement to check and return error, AND explicit requirement NOT to check anything else (added after 2009).	MAY	MAY	N/A	MAY	MAY	SHALL eHx	- TF 2b: 3.41.4.1.3 - TF 3: Table 4.2.4.1-2	- TF 2b: 3.42.4.1.4 - TF 3: Table 4.1-11	
Receiver returns XDSMissingDocumentMetadata if MIME part with Content-Id header not found in metadata	Question for ITI Tech Committee.	MAY	MAY	N/A	MAY	MAY	SHALL eHx	- TF 2b: 3.41.4.1.3 - TF 3: Table 4.2.4.1-2	- TF 2b: 3.42.4.1.4 - TF 3: Table 4.1-11	
Receiver returns InvalidDocumentContent if document content does not match metadata	SHALL	MAY	MAY	N/A	N/A	N/A	MAY	- TF 2b: 3.41.4.1.3	N/A	
Validating: other metadata										
Receiver returns error if HCID of sending system (sourceId) not permitted	MAY, but error code not specified	MAY, but error code not specified	MAY	N/A	MAY	MAY	MAY IHE	- TF 2b: 3.41.4.1.3	- TF 2b: 3.41.6.2	If we have systems that use this whitelisting, they will need a way to know about larger group of senders from QTF. Same issue in other direction. For now will assume if this is done, it will be by checking the directory, and assuming eHx/QTF directories will be cross-pollinated.
Receiver returns UnresolvedReferenceException if UUID in request can't be resolved	SHALL	MAY	MAY	N/A	N/A	N/A	SHALL IHE	- TF 2b: 3.42.4.1.3.5 - TF 3: Table 4.2.4.1-2	- TF 2b: 3.42.4.1.4 - TF 3: Table 4.1-11	
Receiver returns XDSDuplicateUniqueIdInRegistry if new SubmissionSet or Folder for existing uniqueId	SHALL	MAY	MAY	N/A	MAY	MAY	MAY-UND	- TF 2b: 3.42.4.1.3.3 - TF 2b: 3.42.4.1.3.3.7 - TF 3: Table 4.2.4.1-2	- TF 2b: 3.42.4.1.4 - TF 3: Tables 4.1-9, 4.1-10, 4.1-11	
Receiver returns XDSPatientIdDoesNotMatch if objects in a submission set have different patients	SHALL	MAY	MAY	N/A	MAY	MAY	SHALL eHx	- TF 2b: 3.42.4.1.3.3.2 - TF 2b: 3.42.4.1.3.5 - TF 3: Table 4.2.4.1-2	- TF 2b: 3.42.4.1.4 - TF 3: Table 4.1-11	
Receiver returns XDSPatientIdDoesNotMatch if: - Document has different patient from folder - Associated documents have different patients	SHALL	MAY	MAY	N/A	MAY	MAY	SHALL eHx	- TF 2b: 3.42.4.1.3.3.2 - TF 2b: 3.42.4.1.3.5 - TF 3: Table 4.2.4.1-2	- TF 2b: 3.42.4.1.4 - TF 3: Table 4.1-11	
Receiver returns XDSRegistryDeprecatedDocumentError - Association referencing a deprecated document.	SHALL	MAY	MAY	N/A	MAY	MAY	SHALL IHE	- TF 2b: 3.42.4.1.3.5 - TF 3: Table 4.2.4.1-2	- TF 2b: 3.42.4.1.4 - TF 3: Table 4.1-11	
Receiver returns XDSRegistryDuplicateUniqueIdInMessage/XDSRepositoryDuplicateUniqueIdInMessage if uniqueId value was found to be used more than once within the submission	SHALL	MAY	MAY	N/A	MAY	MAY	SHALL eHx	- TF 2b: 3.42.4.1.3.3 - TF 3: Table 4.2.4.1-2	- TF 2b: 3.42.4.1.4 - TF 3: Table 4.1-11	
Receiver returns XDSReplaceFailed: Error detected by the Document Registry during a document replacement	Deprecated	Deprecated	Deprecated	N/A	MAY	MAY	N/A	N/A - deprecated	- TF 2b: 3.42.4.1.4 - TF 3: 4.1.11	
Receiver returns XDSRegistryMetadataError/XDSRepositoryMetadataError if coded values not in defined value sets	SHALL if XDS Affinity Domain constrains	MAY	MAY	N/A	SHOULD (see wiki discussion http://sequoiatechwg.editme.com/Topic-1537135046658)	MAY	SHOULD eHx and use severity of Warning	- TF 2b: 3.42.4.1.3.3.3 - TF 3: Table 4.2.4.1-2	- TF 2b: 3.42.4.1.4 - TF 3: Table 4.1-8 - TF 3: Table 4.1-11	Need to know if QTF is disallowing metadata checks or specifying value sets / mime types
Receiver returns XDSRegistryMetadataError/XDSRepositoryMetadataError if mime type not in allowed set	SHALL if XDS Affinity Domain constrains	MAY	MAY	N/A	MAY	MAY	TBD - Discuss	- TF 2b: 3.42.4.1.3.3.4 - TF 3: Table 4.2.4.1-2	- TF 2b: 3.42.4.1.4 - TF 3: Table 4.1-8 - TF 3: Table 4.1-11	

Feature	XDS.b 2019	XDR 2019	XCDR 2019	XDM 2019	eHx Doc Sub 2011	XDR 2009 (no Vol 2, so mostly equal to XDS.b)	eHx Doc Sub 2020	ITI 2019 refs (XDM omitted)	eHx Doc Sub 2011, ITI 2009 refs (if different)	Notes on XCDR compatibility with QTF
Receiver returns XDSRegistryMetadataError/XDSRepositoryMetadataError if service start time > stop time	SHALL	MAY	MAY	N/A	MAY	MAY	SHALL eHx	- TF 2b: 3.42.4.1.3.3.6 - TF 3: Table 4.2.4.1-2	- TF 2b: 3.42.4.1.4 - TF 3: Table 4.1-8 - TF 3: Table 4.1-11	
Receiver returns XDSRegistryMetadataError/XDSRepositoryMetadataError if adding Document Entry to Folder and both are not Approved	SHALL	MAY	MAY	N/A	MAY	MAY	MAY-UND	- TF 2b: 3.42.4.1.3.4 - TF 3: Table 4.2.4.1-2	- TF 2b: 3.42.4.1.4 - TF 3: 4.1.11	
Receiver returns XDSRegistryMetadataError/XDSRepositoryMetadataError if IsSnapshotOf Association does not relate a Stable to On-Demand entry	SHALL	MAY	MAY	N/A	N/A	N/A	MAY-UND	- TF 2b: 3.42.4.1.3.5 - TF 3: Table 4.2.4.1-2	- TF 2b: 3.42.4.1.4 - TF 3: Table 4.1-11	
Receiver returns XDSRegistryMetadataError/XDSRepositoryMetadataError if missing required metadata field	SHALL	MAY	MAY	N/A	MAY	MAY	SHALL eHx	- TF 2b: 3.42.4.1.3.5 - TF 3: Table 4.2.4.1-2	- TF 2b: 3.42.4.1.4 - TF 3: Table 4.1-11	
Receiver returns XDSRegistryMetadataError/XDSRepositoryMetadataError if any remaining "shalls" in TF 3: 4 are violated in metadata. Example: 4.2.1.3, folder nested inside another folder.	SHALL implied. These cover the structure and constraints of the metadata, but not all are explicitly required to be checked. Question for ITI Tech Committee.	MAY	MAY	N/A	MAY	MAY	SHOULD eHx	- TF 2b: 3.42.4.1.3.3 - TF 3: 4.1	- TF 2b: 3.42.4.1.4 - TF 3: Table 4.1-11	
General/unspecified errors										
Receiver returns XDSRegistryBusy/XDSRepositoryBusy if too much activity	MAY	MAY	MAY	N/A	MAY	MAY	MAY eHx	- TF 2b: 3.41.4.1.3 - TF 3: Table 4.2.4.1-2	- TF 2b: 3.42.4.1.4 - TF 3: Table 4.1-11	
Receiver returns XDSRegistryError/XDSRepositoryError if internal error	MAY	MAY	MAY	N/A	MAY	MAY	MAY eHx	- TF 2b: 3.41.4.1.3 - TF 3: Table 4.2.4.1-2	- TF 2b: 3.42.4.1.4 - TF 3: Table 4.1-11	
Receiver returns XDSRegistryNotAvailable if Repository was unable to access the Registry	MAY	MAY	MAY	N/A	MAY	MAY	MAY eHx	- TF 2b: 3.41.4.1.3 - TF 3: Table 4.2.4.1-2	- TF 2b: 3.42.4.1.4 - TF 3: Table 4.1-11	
Receiver returns XDSRegistryOutOfResources/XDSRepositoryOutOfResources if resources are low	MAY	MAY	MAY	N/A	MAY	MAY	MAY eHx	- TF 2b: 3.41.4.1.3 - TF 3: Table 4.2.4.1-2	- TF 2b: 3.42.4.1.4 - TF 3: Table 4.1-11	