



pan-Canadian FHIR Exchange (CA:FeX) Interoperability Specifications

Version: 0.2

Type: Draft

Release Date: March 9, 2022

Table of Contents

| | | |
|----------|--|-----------|
| 1 | CA:FeX Interoperability Specifications..... | 5 |
| 1.1 | Introduction | 5 |
| 1.2 | Intended Audience..... | 5 |
| 1.3 | Purpose..... | 5 |
| 1.4 | Glossary of Terms and Acronyms..... | 6 |
| 1.5 | Preface..... | 15 |
| 1.5.1 | Context | 16 |
| 1.5.2 | Introduction to IHE | 16 |
| 1.5.3 | How to Read This Document..... | 16 |
| 1.5.4 | Document Conventions | 17 |
| 1.5.5 | Requirements Language..... | 18 |
| 1.5.6 | Methodology | 18 |
| 1.5.7 | Introduction to a Use-Case Driven Approach | 18 |
| 1.5.8 | Release Cycle..... | 18 |
| 1.6 | CA:FeX Use Case Overview..... | 19 |
| 1.6.1 | Use Case Overview | 19 |
| 1.6.2 | In-Scope..... | 19 |
| 1.6.3 | Use Case Actors and Services | 19 |
| 1.6.4 | Design Constraints & Assumptions | 20 |
| 1.7 | Core Interoperability Specification Requirements | 21 |
| 1.7.1 | Actor Mapping to Interoperability Specification | 21 |
| 1.7.2 | Table 1 Interoperability Conformance Requirements for Use Case 1: Create and Submit Document..... | 22 |
| 1.7.3 | Table 2 Interoperability Conformance Requirements for Use Case 2: Query and Retrieve Document..... | 23 |
| 1.8 | Actor Conformance | 25 |
| 1.8.1 | Constraints on Use Case Actors..... | 25 |
| 1.9 | Privacy & Security Guidance..... | 29 |
| 1.9.1 | Privacy Considerations | 29 |
| 1.9.2 | Security Considerations..... | 29 |
| 1.10 | CA:FeX Actors and Transactions..... | 30 |
| 2 | Use Cases and Definitions | 31 |
| 2.1 | Use Case Index | 31 |

| | | |
|----------|--|-----------|
| 2.2 | Requirement Priority Definitions | 31 |
| 2.3 | UC-01 Create and Submit Document | 33 |
| 2.4 | UC-02 Query and Retrieve Document | 37 |
| 3 | Exchanging Documents in FHIR | 41 |
| 3.1 | Overview | 41 |
| 3.2 | FHIR Version | 42 |
| 3.3 | Submit a Document | 43 |
| 3.3.1 | Out of Scope | 43 |
| 3.3.2 | Use Cases | 43 |
| 3.3.3 | CA:FeX Transactions | 43 |
| 3.3.4 | HTTP Operations | 44 |
| 3.3.5 | Document Submit Patterns | 44 |
| 3.3.6 | Document Submit Patterns: Augmenting Base Capabilities | 45 |
| 3.3.7 | Cutting Edge Capabilities | 45 |
| 3.4 | Search for a Document | 47 |
| 3.4.1 | Scope | 47 |
| 3.4.2 | Use Cases | 47 |
| 3.4.3 | CA:FeX Transactions | 47 |
| 3.4.4 | HTTP Operations | 47 |
| 3.4.5 | Considerations for Document Search Patterns | 48 |
| 3.4.6 | Document Search Patterns: CA:FeX-2A (Search Against FHIR Assembled Documents Repository) | 48 |
| 3.4.7 | Document Search Patterns: CA:FeX-2B (Search Against Hybrid Documents Repository) | 50 |
| 3.4.8 | Cutting Edge Capabilities | 52 |
| 3.5 | Retrieve a Document | 53 |
| 3.5.1 | Use Cases | 53 |
| 3.5.2 | CA:FeX Transactions | 53 |
| 3.5.3 | HTTP Operations | 53 |
| 3.5.4 | Considerations for Document Retrieve Patterns | 53 |
| 3.5.5 | Document Retrieve Patterns | 53 |
| 3.5.6 | Document Retrieve Patterns: Supported Retrieval Parameters | 54 |
| 3.5.7 | Cutting Edge Capabilities | 54 |
| 3.6 | Response Handling | 55 |
| 3.6.1 | HTTP Status Codes | 55 |

| | |
|---|-----------|
| 3.6.2 Operation Outcome | 56 |
| 4 Reference Architecture | 58 |
| 4.1 Purpose..... | 58 |
| 4.2 Intended Audience..... | 58 |
| 4.3 Overview..... | 58 |
| 4.4 How To Read The Reference Architecture | 58 |
| 4.5 How To Use The Reference Architecture | 59 |
| 4.6 High Level View - Reference Architecture for CA:FeX | 59 |
| 4.7 Foundational IHE Profiles..... | 61 |
| 4.8 Sequence Diagram for UC-01: Create and Submit Document..... | 64 |
| 4.9 Sequence Diagram for UC:02: Query and Retrieve Document | 66 |
| 4.9.1 Additional Considerations..... | 66 |

∞ **Viewing large images in this file:**

The CA:FeX Specifications contain several large tables and diagrams (e.g., use cases diagrams, sequence diagrams, etc.). In some cases, for optimal viewing, it is recommend to view these in the online specifications available [here](#) under Published Content.

1 CA:FeX Interoperability Specifications

1.1 Introduction

Interoperability enables information to flow seamlessly between different solutions and devices. When different parts of the health care system are interoperable with each other, they can “speak the same language.” Interoperability improves continuity of care, collaboration between health care providers and patient access to their health information. By breaking down data silos, it also reduces inefficiencies and redundancies within the health care system.

Connection, collaboration and communication have never been more important for the health care system. Increased use of digital health solutions within healthcare has highlighted the need for safe and efficient electronic sharing of information across the circle of care. Continuing to improve Canadian health care will necessitate work in interoperability — connected systems are healthier systems.

In support of the provinces and territories, Canada Health Infoway (Infoway) is facilitating a national collaborative effort to advance interoperability. While there are many interoperability-related challenges, this specification addresses standardized sharing of vital patient information for the benefit of health care providers and patients using FHIR based information exchange. This FHIR based information exchange is similar to and accomplishes the same objectives as a Health Information Exchange (HIE).

The [Office of the National Coordinator for Health Information Technology](#) (ONC) defines HIE as:

Electronic health information exchange (HIE) allows doctors, nurses, pharmacists, other health care providers and patients to appropriately access and securely share a patient’s vital medical information electronically—improving the speed, quality, safety and cost of patient care.

While electronic health information exchange cannot replace provider-patient communication, it can greatly improve the completeness of patient’s records, (which can have a big effect on care), as past history, current medications and other information is jointly reviewed during visits.

Appropriate, timely sharing of vital patient information can better inform decision making at the point of care and allow providers to avoid readmissions, avoid medication errors, improve diagnoses and decrease duplicate testing.

1.2 Intended Audience

The intended audience of the CA:FeX Interoperability Specifications (Canadian FHIR Exchange (CA:FeX)), includes but is not limited to:

- Those interested in integrating healthcare information systems and workflows;
- IT departments of healthcare institutions;
- Technical staff of clinical solution vendors;
- Experts involved in standards development; and
- Software developers.

1.3 Purpose

The purpose of this document is to address the following functionality for release 1:

- Address two FHIR HIE use cases;
- Provide a detailed set of requirements (including actors, transactions and references to specific profiles and standards);
- Describe the implementation patterns that enable the exchange of clinical information using CA:FeX; and
- Describe the set of requirements that complement the set of IHE Profiles and HL7 FHIR® Profiles required by the CA:FeX specifications with Canadian specific constraints.

1.4 Glossary of Terms and Acronyms

The following table provides a list of terms and acronyms that you may encounter throughout the pan-Canadian interoperability specifications (e.g. PS-CA, CA:FeX) and/or in the prototyping and validation information.

| Term / Acronym | Meaning |
|--|---|
| ATNA | The Audit Trail and Node Authentication (ATNA) Profile specifies the foundational elements needed by all forms of secure systems: node authentication, user authentication, event logging (audit), and telecommunications encryption. It is also used to indicate that other internal security properties such as access control, configuration control, and privilege restrictions are provided. (Source: https://profiles.ihe.net/ITI/TF/Volume1/ch-9.html) |
| Author (e.g., PS-CA Author) | A health care provider who authors and/or curates clinical data (e.g. Patient Summary). |
| Business/Legal Interoperability Requirements | Requirements that enable independent organizations to execute a collaborative process or service. |
| Business Requirements: Non-Testable | Business requirements that are not directly traceable to an IHE profile in the PS-CA specifications (e.g., provided for consideration and to support and provide guidance to implementers of the PS-CA). |
| Business Requirements: Testable | Business requirements that are directly traceable to an IHE profile in the PS-CA specifications. |
| CA:FeX | The CA:FeX Interoperability Specifications (Canadian FHIR Exchange (CA:FeX)) seek to promote FHIR RESTful exchange patterns, developed by industry-leading FHIR standards that can be applied on top of an existing non-FHIR infrastructure just as easily as it can be applied on top of FHIR servers. |
| CCDD | The Canadian Clinical Drug Data Set (CCDD) is the drug terminology for use in digital health solutions such as electronic prescribing in Canada. |
| Central Infrastructure | A Central Infrastructure collects health information from participating organizations and stores the information in a centralized place. The Infrastructure also provides access control. Typically, the central infrastructure is under jurisdictional control. |

| Term / Acronym | Meaning |
|---------------------------------|---|
| Clinical Data eXchange (CDX) | <p>CDX is a clinical distribution service developed by Interior Health. Northern Health (NH) and Interior Health (IH) have collaborated to facilitate the sharing of Health Authority clinical information to participating provider EMR systems using this service.</p> <p>(Sources: https://infocentral.infoway-inforoute.ca/en/resources/docs/coordofcare/1406-clinical-document-exchange-bc-cdx-technical-overview-coc-sep27-16 https://www.intrahealth.com/sites/default/files/docs/Clinical-Data-eXchange-communication-from-Intrahealth-and-CDX-Team.pdf)</p> |
| Clinical Data Repository | A clinical data repository is built around the HL7 FHIR standard used for storing clinical data. |
| Clinical Solution | Any combination of health information technology assets and processes that enables clinical data to be communicated, managed and dispositioned between a Producer and a Consumer. Clinical Solutions can be comprised of various Producer and Consumer systems including: EMR, HIS, CIS, PHR, EHR or any combination of these systems. |
| Conformance Testing | Conformance testing is a formal process of assessment focused on ensuring clinical solutions and systems accurately implement a particular specification (e.g. PS-CA Specifications) by ensuring there is conformance to the stated parameters that are being claimed in the standard. |
| Consumer (e.g., PS-CA Consumer) | A health records system (e.g., EMR, HIS, CIS, PHR, Patient Portal or EHR) that enables access to or receipt of a clinical document (e.g. Patient Summary-CA) by an authorized health care provider or the subject of care/patient. |
| Cross Border, Scheduled Care | Scheduled care of a resident of Canada that is delivered in/by another country. |
| Cross Border, Unscheduled Care | Unscheduled care of a resident of Canada that is delivered in/by another country. |
| CT | <p>The Consistent Time Integration Profile (CT) provides a means to ensure that the system clocks and time stamps of the many computers in a network are well synchronized. This profile specifies synchronization with a median error less than 1 second. This is sufficient for most purposes.</p> <p>(Source: https://profiles.ihe.net/ITI/TF/Volume1/ch-7.html)</p> |
| DIN | A Drug Identification Number (DIN) is a computer-generated eight digit number assigned by Health Canada to a drug product prior to being marketed in Canada. |

| Term / Acronym | Meaning |
|--|--|
| Document Repository (Local or Central) | A document repository is a shared storage space for clinical documents (Patient Summaries) that can be hosted locally (i.e., at the document producer) or at the HIE Central Infrastructure and can be accessed by authorized users. |
| DPD | The Drug Product Database (DPD) is used to find drugs authorized for sale by Health Canada. The DPD is updated nightly and includes: availability of the drug in Canada. |
| Electronic Health Record (EHR) | <p>The EHR represents the Clinical Solution that contains a secure and private collection of a patient's health information in a digital format, which is shareable across different health care settings / clinical solutions that are integrated. The EHR facilitates better sharing and interpretation of health information among the health care professionals involved in the care of the patient. For example:</p> <ul style="list-style-type: none"> • CareConnect is British Columbia's secure, view-only EHR solution. It offers healthcare providers access to an integrated, provincial view of patient-centric information available 24/7 to support the delivery of patient care. • HEALTHe NL is the Newfoundland & Labrador provincial EHR. HEALTHe NL will provide more accurate and reliable data to support improved health care delivery, decision-making and policy and create improved accountability, stability and efficiency in the provincial health care system. • Netcare is Alberta's name for all the projects related to the provincial EHR - a secure and confidential electronic system of Alberta patients' health information: a single, comprehensive, and integrated patient record. • Other clinical systems: In some health authorities, other clinical systems may act as an EHR, holding the patient summary information. |
| Extensible PS-CA Dataset | <p>Extensible PS-CA Dataset: PS-CA content that can be extended for use in a PS-CA use case scenario that complements the primary PS-CA use cases.</p> <p>*Note: Extensible PS-CA Dataset refers to the addition of data domains such as Family History.</p> |
| FHIR® | <p>HL7 Fast Healthcare Interoperability Resources (FHIR), is expected to be a next generation standards framework created by HL7. FHIR combines the best features of HL7's Version 2, Version 3 and product lines while leveraging the latest web standards and applying a tight focus on implementability.</p> <p>(Source: http://www.hl7.org/implement/standards/fhir/)</p> |
| FHIR Repository | A FHIR repository is a clinical data repository built around the HL7 FHIR standard used for storing clinical data. |

| Term / Acronym | Meaning |
|--|--|
| Gazelle | Gazelle is a suite of virtual tools, developed by IHE Europe used to support interoperability testing. Gazelle will allow jurisdictions and vendors an opportunity to validate the role they will be playing in an ecosystem and ensure they are able to satisfy the interoperability requirements. Gazelle offers several self-serve, self-test and innovation opportunities for jurisdictions and vendors to test their alignment to the represented integration profiles. |
| HCP | Health Care Provider |
| Health Information Access Layer (HIAL) | An interface specification for the EHR infostructure that defines service components, service roles, information model and messaging standards required for the exchange of EHR data and execution of interoperability profiles between EHR services. (Source: https://www.infoway-inforoute.ca/en/component/edocman/resources/technical-documents/391-ehrs-blueprint-v2-full ; Page.340) |
| Health Information Exchange (HIE) | Electronic health information exchange (HIE) allows doctors, nurses, pharmacists, other health care providers and patients to appropriately access and securely share a patient's vital medical information electronically—improving the speed, quality, safety and cost of patient care. While electronic health information exchange cannot replace provider-patient communication, it can greatly improve the completeness of patients' records, (which can have a big effect on care), as past history, current medications and other information is jointly reviewed during visits. Appropriate, timely sharing of vital patient information can better inform decision making at the point of care and allow providers to avoid readmissions, avoid medication errors, improve diagnoses and decrease duplicate testing. (Source: https://www.healthit.gov/topic/health-it-and-health-information-exchange-basics/what-hie) |
| Health Records System | A health records system may include an electronic medical records system, a hospital information system, a clinical information system, an electronic health records system or a personal health records system. The term is broadly used to describe system actors that may produce and/or consume a PS-CA. Jurisdictional implementation patterns will determine which systems are used to create, access, consume and manage patient summaries. |
| HIS | Health Information System |
| Health Level 7 (HL7) | Founded in 1987, HL7 is a not-for-profit standards developing organization dedicated to providing a comprehensive framework and related standards for the exchange, integration, sharing, and retrieval of electronic health information that supports clinical practice and the management, delivery and evaluation of health services. (Source: http://www.hl7.org/about/index.cfm?ref=nav) |

| Term / Acronym | Meaning |
|---|---|
| HL7 Fast Healthcare Interoperability Resources (FHIR) | <p>Expected to be a next generation standards framework created by HL7. FHIR combines the best features of HL7's Version 2, Version 3 and product lines while leveraging the latest web standards and applying a tight focus on implementability.</p> <p>(Source: http://www.hl7.org/implement/standards/fhir/)</p> |
| Health Records System | <p>A health records system may include an electronic medical records system, a hospital information system, a clinical information system, an electronic health records system or a personal health records system. The term is broadly used to describe system actors that may produce and/or consume a PS-CA. Jurisdictional implementation patterns will determine which systems are used to create, access, consume and manage patient summaries.</p> |
| Information/Semantic Interoperability Requirements | <p>Requirements for syntax and semantics such that data exchanged between health record systems can be interpreted and the meaning of the data ascertained.</p> |
| Integrating the Healthcare Enterprise (IHE) | <p>IHE is an initiative by healthcare professionals and industry to improve the way computer systems in healthcare share information. IHE promotes the coordinated use of established standards such as DICOM and HL7 to address specific clinical needs in support of optimal patient care. Systems developed in accordance with IHE communicate with one another better, are easier to implement, and enable care providers to use information more effectively.</p> <p>(Source: https://www.ihe.net/)</p> |
| IHE Actor | <p>IHE Actors are responsible for producing, managing and/or acting on information in the context of an IHE Profile (e.g., Primary Care Provider, EMR, EHR, etc.).</p> <p>(Source: https://wiki.ihe.net/index.php/Actors)</p> |
| IHE Domain | <p>IHE Domains are responsible for the development and maintenance of the IHE Technical Frameworks that document the Integration Profiles. Each Domain manages Integration Profiles in a particular part of healthcare (e.g., Virtual Care).</p> <p>(Source: https://wiki.ihe.net/index.php/Domains)</p> |
| IHE Profiles | <p>IHE Profiles describe specific solutions to interoperability problems. Profiles specify how "Actors" use standards to address a specific healthcare use case (e.g., Medication, Allergy Intolerance, etc.).</p> <p>(Source: https://wiki.ihe.net/index.php/Profiles)</p> |

| Term / Acronym | Meaning |
|---------------------------------------|---|
| IHE Transactions | <p>IHE Transactions are interactions between actors that communicate the required information through standards-based messages (e.g., patient look-up query, send patient summary information, etc.).</p> <p>(Source: https://wiki.ihe.net/index.php/PCC_TF-1/About)</p> |
| International Patient Summary (IPS) | <p>The IPS is a a minimal, non-exhaustive set of data elements defined by ISO/EN 17269 and realized by HL7 in both CDA and FHIR. The IPS is a snapshot clinical document that can be used for planned or unplanned care of a person locally or across borders. It emphasizes the data required and the necessary conformance of the use cases for an international patient summary.</p> <p>(Source: https://wiki.ihe.net/index.php/International_Patient_Summary_(IPS))</p> |
| Interoperability | <p>Interoperability enables information to flow seamlessly between different solutions and devices. When different parts of the health system are interoperable with each other, they can “speak the same language.” Interoperability improves continuity of care, collaboration between health providers and patient access to their health information. By breaking down data silos, it also reduces inefficiencies and redundancies within the health system.</p> <p>Connection, collaboration and communication have never been more important for the health system. Increased use of virtual care has highlighted the need for safe and efficient electronic sharing of information across the circle of care. Continuing to improve Canadian health care will necessitate work in interoperability — connected systems are healthier systems.</p> <p>For more information about interoperability, please visit Canada Health Infoway - Interoperability.</p> |
| IUA | <p>The Internet User Authorization (IUA) Profile provides support for authorizing network transactions when using HTTP RESTful transports. IHE has authorization profiles for the Web Services and SOAP based transactions, and this profile provides an authorization profile for the HTTP RESTful transactions.</p> <p>(Source: https://profiles.ihe.net/ITI/TF/Volume1/ch-34.html)</p> |
| Local, Scheduled Care | <p>Scheduled care of a resident of Canada that is delivered in/by the Canadian health care system. This includes care provided in federal, provincial and territorial jurisdictions, as well as cross-jurisdictional care.</p> |
| Local, Unscheduled Care | <p>Unscheduled care of a resident of Canada that is delivered in/by the Canadian health care system. This includes care provided in federal, provincial and territorial jurisdictions, as well as cross-jurisdictional care.</p> |
| Longitudinal Electronic Health Record | <p>A longitudinal electronic health record is a single comprehensive patient record comprised of data from numerous data sources across the healthcare continuum.</p> |

| Term / Acronym | Meaning |
|----------------------------|---|
| MHD | <p>The Mobile access to Health Documents (MHD) Profile defines one standardized interface to health document sharing (a.k.a. an Application Programming Interface (API)) for use by mobile devices so that deployment of mobile applications is more consistent and reusable.</p> <p>(Source: https://profiles.ihe.net/ITI/MHD/index.html)</p> |
| Netcare | <p>Alberta Netcare is the name for all the projects related to the provincial Electronic Health Record (EHR) - a secure and confidential electronic system of Alberta patients' health information: a single, comprehensive, and integrated patient record.</p> <p>The EHR stores information about:</p> <ul style="list-style-type: none"> • laboratory tests, • dispenses of pharmaceuticals (drugs), • hospital discharge reports, and • diagnostic imaging |
| Patient Portal | <p>A patient portal is a web-based access point that enables secure patient access to personal health information and other self-serve health IT services. For example, a patient portal can be hosted on an EMR solution.</p> |
| Patient Proxy | <p>An individual or entity that has the authority to act on behalf of a subject of care/ patient. Proxies can include parents of dependent children, parents of dependent adults, powers of attorney etc.</p> |
| Patient Summary-CA (PS-CA) | <p>An electronic patient summary for use at the point of care comprised of, at minimum, the required elements of the Patient Summary-CA data set and specifications. The PS-CA is a health record extract, at a snapshot in time, comprised of a standardized collection of clinical and contextual information (retrospective, concurrent, prospective), including the minimum necessary and sufficient data to inform a patient's treatment at the point of care. The PS-CA is condition-independent and specialty-agnostic, irrespective of the condition of the patient or the treatment sought or specialty of the provider delivering care.</p> |
| PDQm | <p>The Patient Demographics Query for Mobile (PDQm) Profile defines a lightweight RESTful interface to a patient demographics supplier leveraging technologies readily available to mobile applications and lightweight browser based applications.</p> <p>(Source: https://profiles.ihe.net/ITI/TF/Volume1/ch-38.html)</p> |
| PIXm | <p>The Patient Identifier Cross-reference for Mobile (PIXm) Profile provides RESTful transactions for mobile and lightweight browser-based applications to create, update and delete patient records in a Patient Identifier Cross-reference Manager and to query the Patient Identifier Cross-reference Manager for a patient's cross-domain identifiers.</p> <p>(Source: https://profiles.ihe.net/ITI/TF/Volume1/ch-41.html)</p> |

| Term / Acronym | Meaning |
|---|---|
| PMIR | <p>The Patient Master Identity Registry (PMIR) Profile supports the creating, updating and deprecating of patient master identity information about a subject of care, as well as subscribing to changes to the patient master identity, using the HL7 FHIR standard resources and RESTful transactions.</p> <p>(Source: https://www.ihe.net/uploadedFiles/Documents/ITI/IHE_ITI_Suppl_PMIR.pdf)</p> |
| Producer (e.g., PS-CA Producer) | <p>A health records system (e.g., EMR, HIS, CIS, PHR, or EHR) that creates/produces a clinical document (e.g. Patient Summary) in response to a request from an authorized health care provider, the subject of care or another authorized health records system.</p> |
| Projectathon | <p>A Projectathon is an important step and a best-practice approach in testing and validation of a specification package, where implementers collaborate to test their solutions using methodology and tools that accelerate interoperability. A Projectathon provides an opportunity for participants to test their systems among themselves and against a reference environment. It is also an opportunity to collaborate among peers to enable hands-on knowledge exchange.</p> |
| PS-CA Solution | <p>Any combination of health information technology assets and processes that enables a Patient Summary-CA to be created, communicated, managed and dispositioned between a PS-CA Producer and a PS-CA Consumer. Patient Summary-CA Solutions can be comprised of various Producer and Consumer systems including: EMR, HIS, CIS, PHR, EHR or any combination of these systems.</p> |
| PS-CA Specifications | <p>pan-Canadian Patient Summary Interoperability Specifications: The pan-Canadian Patient Summary Interoperability Specification is an implementable, testable specification, based on the IHE International Patient Summary specification and the HL7 IPS Implementation Guide. For more information on the PS-CA Specifications, please go here.</p> |
| PT | Provinces and Territories |
| SUT | System Under Test |
| Technical Interoperability Requirements | <p>Requirements for one health record system to send data to another health record system and for the receiving system to acknowledge receipt of the data payload.</p> |
| Terminology | <p>Collection of uniquely identifiable concepts with associated representations, designations, associations and meanings.</p> |

| Term / Acronym | Meaning |
|----------------|--|
| XDM | <p>Cross-Enterprise Document Media Interchange (XDM) provides document interchange using a common file and directory structure over several standard media types. This permits the patient to use physical media to carry medical documents. This also permits the use of person-to-person email to convey medical documents. XDM supports the transfer of data about multiple patients within one data exchange.</p> <p>(Source: https://profiles.ihe.net/ITI/TF/Volume1/ch-16.html)</p> |
| XDS | <p>The Cross-Enterprise Document Sharing (XDS) IHE Integration Profile facilitates the registration, distribution and access across health enterprises of patient electronic health records.</p> <p>(Source: https://profiles.ihe.net/ITI/TF/Volume1/ch-10.html)</p> |

1.5 Preface

The CA:FeX Interoperability Specifications (Canadian FHIR Exchange (CA:FeX)) seek to promote FHIR RESTful exchange patterns, developed by industry-leading FHIR standards that can be applied on top of an existing infrastructure just as easily as it can be applied on top of FHIR servers.

The current version of CA:FeX is focused on the FHIR RESTful exchange of documents, which FHIR offers multiple structures and patterns to achieve. CA:FeX aims to provide clarity to implementers by identifying some of the choices currently available using FHIR, ranging from simple to a higher level of sophistication. As the development of CA:FeX continues, the Interoperability Specifications may evolve into a more formal Integration Profile (similar to existing international profiles reviewed below) that provides more comprehensive guidance on FHIR RESTful exchange patterns than what currently exists today.

The scope of CA:FeX may evolve beyond document-based exchange to explore a more atomic nature of health information exchange in the future. This will depend on the development of a multi-phased CA:FeX roadmap, the evolving needs of the Canadian health care market, and emerging trends within the international FHIR health information exchange community.

Upon review of legacy IHE profiles, existing guidance is either based on the enablement of document sharing within non-FHIR infrastructures or through the narrow use of FHIR resources. Given the evolving needs of the Canadian market, these legacy profiles may not be entirely sufficient to enable the FHIR RESTful exchange of documents. Two key existing IHE profiles that were contemplated, including some of their limitations, are listed below:

- *Cross-Enterprise Document Sharing (XDS)*: This IHE profile is focused on providing a standards-based specification for the sharing of documents which is limited to non-FHIR infrastructures.
- *Mobile access to Health Documents (MHD)*: This IHE profile is designed for the utilization of FHIR standard to communicate and exchange documents, including, as an option, acting as a proxy to systems that use XDS. MHD leverages FHIR resources (i.e., DocumentReference) as a standard method for clients to find a document. To ensure this IHE profile can be used without regard for how information is stored (e.g., XDS infrastructure, FHIR, or another storage system) the profile applies constraints to the FHIR DocumentReference resource. The current limitation of MHD, however, requires that a document is submitted in FHIR Binary format and does not include the use of a FHIR Bundle (of type Document) to represent documents. Further, MHD entails a multi-step document retrieval process (list/find step followed by the retrieval step) which may not be the single approach implementers will take going forward.

Through a market scan, it was observed that in line with MHD, existing RESTful driven FHIR Implementation Guides that include document exchange (e.g., US Core, IPA, PACIO, IPS, etc.) have begun utilizing a pattern, leveraging the DocumentReference Resource. These Implementation Guides, however, enable querying in a single way and return pointers to document content, wherever it is stored and irrespective of the format (e.g., binary or FHIR-assembled). FHIR Search Parameters and FHIR Operations have been developed to augment the capabilities of this pattern to more easily get back what was requested and enable the offering of documents in the expected format without having to change the underlying data model / document and lifecycle practices.

Within the context of the Canadian market, there is an indication that the level of readiness for exchange patterns outlined by international FHIR Implementation Guides, such as the US Core, may not be appropriate for implementers in the current state. As such, simpler exchange patterns are also a viable alternative. One such alternative is narrowing the focus to the exchange of non-binary documents and utilizing the appropriate FHIR Search Parameters and FHIR Operations based on the Composition Resource to search and retrieve documents.

Given the preliminary findings, both internationally and domestically, CA:FeX outlines the use of FHIR Search and FHIR Operations based on both (i) the DocumentReference Resource and (ii) the Composition Resource. While these methods are surfaced by CA:FeX to drive adoption and feedback, this is still being tested and addressed by the FHIR document exchange community. Furthermore, the current methods explored by this version of CA:FeX do not represent the only FHIR methods available to enable document exchange or, more broadly, health information

exchange as a whole. As international specifications mature and the Canadian market evolves, CA:FeX will dig deeper into the more sophisticated alternatives for health information exchange in order to guide the Canadian market in the same direction.

Due to the continuous development of the market, there will be an opportunity to encompass a broader set of use cases that may deal with the exchange of more atomic information going forward.

In parallel with CA:FeX, Canada Health Infoway is also facilitating a national collaborative effort to develop the [pan-Canadian Patient Summary Interoperability Specifications \(PS-CA\)](#), an implementable, testable specification, based on the IHE International Patient Summary specification and the HL7 IPS Implementation Guide. The PS-CA references the CA:FeX Interoperability Specifications as an optional implementation pattern for submitting, searching and retrieving a Patient Summary document.

1.5.1 Context

The CA:FeX Interoperability Specifications are published to a public space within Canada Health Infoway's InfoScribe and are also available as a downloadable document, [here](#). InfoScribe is a web-based tool developed for jurisdictions and vendors to create, publish, and collaborate on clinical requirements and specifications for interoperability solutions. Teams can document, share, and discuss content, files, ideas, specs, mock-ups, diagrams, and projects. A link to the online published content and the downloadable documentation will be published with each release of CA:FeX.

1.5.2 Introduction to IHE

Integrating the Healthcare Enterprise (IHE) is an international initiative to promote the use of standards to achieve interoperability among health information technology (HIT) systems and effective use of electronic health records (EHRs). IHE provides a forum for care providers, HIT experts and other stakeholders in several clinical and operational domains to reach consensus on standards-based solutions to critical interoperability issues.

The primary output of IHE is system implementation guides, called IHE profiles. IHE publishes each profile through a well-defined process of public review and Trial Implementation and gathers profiles that have reached Final Text status into an IHE Technical Framework.

1.5.3 How to Read This Document

This document contains the following sections:

| Document Section | Description | Target Audience |
|---|--|---------------------|
| CA:FeX Interoperability Specifications | CA:FeX is an implementable, testable specification based on HL7 FHIR Implementation Guides. It defines building blocks to enable FHIR Health Information Exchange (HIE) implementation patterns. CA:FeX building blocks are configurable to address necessary Canadian jurisdictional variances. The CA:FeX Interoperability Specifications, written in line with international best practices, contain the information necessary for an implementer to consume and develop the components necessary for creating, consuming and sharing clinical data and may be applied to existing and new information systems. | Solution Developers |

| Document Section | Description | Target Audience |
|-------------------------------------|--|------------------------------------|
| Use Cases & Definitions | The Use Cases & Definitions present the broader context for clinical, business, interoperability and solution development considerations that were discovered during the development of the CA:FeX Interoperability Specifications. This section defines the healthcare problem that CA:FeX addresses and includes healthcare use cases and interoperability requirements in terms that will be traceable to the content in the Reference Architecture, which defines the actors and their interactions in scope for the CA:FeX Interoperability Specifications. | CTOs, CMIOs, CIOs, PTs and Vendors |
| Exchanging Documents in FHIR | The Exchanging Documents in FHIR section contains methods for implementing the CA:FeX transactions specifically for document exchange using FHIR RESTful APIs. As the content is refined and evolved it will grow into a FHIR implementation guide. | Solution Developers |
| Reference Architecture | The Reference Architecture contains background information on the abstracted CA:FeX actors and transactions for stakeholders who are not familiar with the IHE methodology. It describes baseline information on the recommended and optional IHE profiles and includes links to the IHE source documentation where stakeholders can get additional details on each CA:FeX actor and transaction. | CTOs, CMIOs, CIOs, PTs and Vendors |

Within the Pan-Canadian CA:FeX Interoperability Specifications, you can expect the following subsections:

- **Introduction & Preface:** Contains an introduction to the CA:FeX Interoperability Specifications. This section contains a summary of the context, document purpose and scope, as well as other content to help orient the first-time reader to the topic of these specifications and how they relate to other specifications in the digital health ecosystem in Canada.
- **Use Case Overview:** Describes the Use Cases, including design constraints and assumptions and the flows of information that will be specified in the CA:FeX Interoperability Specifications. This section also references scenarios that describe how the specified flows may be used in the Canadian context.
- **Core Interoperability Specification Requirements:** Establishes the Core Interoperability Requirements for the CA:FeX Interoperability Specifications with respect to a FHIR Health Information Exchange (HIE) implementation pattern. This section also provides mapping of use case actors to the technical actors of the CA:FeX Interoperability Specifications and the services they are supporting, which are aligned with the flow captured in the sequence diagrams included in the Reference Architecture section.
- **CA:FeX Actor Conformance:** Establishes the Actor Conformance Requirements for the CA:FeX Interoperability Specifications.
- **Privacy & Security Guidance:** Provides a reference to the upcoming Privacy Toolkit to be released in April, Interoperability Challenge Accepted: A Privacy Toolkit for Canadian Jurisdictions, that will address privacy considerations at a pan-Canadian level and describes key Security considerations for the CA:FeX Interoperability Specifications.
- **CA:FeX Actors and Transactions:** Illustrates the abstracted actors and transactions in scope for the CA:FeX Interoperability Specifications.

1.5.4 Document Conventions

The CA:FeX Interoperability Specifications will be numbered according to this format:

- **Name + Version + Stage**, where name refers to the name of the document, version refers to the versioning history of the document and stage refers to its stage in implementation such as “Draft For Use” or “DU”.
- Key documents will evolve during review cycles from version 0.x to v1.0.

1.5.5 Requirements Language

The following conventions are used to specify requirement levels for the business requirements of the CA:FeX Interoperability Specifications:

- **Shall**: used to indicate a **required** requirement.
- **Should**: used to indicate that a requirement is **recommended** and should be considered as best practice for implementation, but not required (i.e., it is optional) for implementation.
- **May**: used to indicate that a requirement is permissible / **optional**, but not required for implementation.
- **Shall not**: used to indicate that an element or action is prohibited.

Additional information on the CA:FeX business requirements can be found in the Use Cases and Definitions section.

1.5.6 Methodology

The CA:FeX Interoperability Specifications have been developed based on international research and stakeholder consultations with HIE Subject Matter Advisors, where this was socialized and validated with participating jurisdictions and vendors through Coordinating Table Meetings, Executive Table Meetings, stakeholder workshops and 1-on-1 meetings to further refine the specifications.

1.5.7 Introduction to a Use-Case Driven Approach

The following use case-driven approach was utilized in the development of the CA:FeX Interoperability Specifications:

- **Baseline**: Develop foundational Use Cases, Use Case Scenarios and Business Requirements for FHIR Health Information Exchange (HIE).
- **Collaborate**: Collaborate with jurisdictions, clinical SMEs, technical SMEs, vendors, participating organizations to develop and refine detailed artefacts.
- **Review**: Review and provide feedback into artefacts through engagement workshops and input gathering.
- **Publish**: Publish artefacts for broader stakeholder consultation.
- **Recommend**: Recommend draft artefacts for approval.
- **Iterate**: Continue to refine as per testing and priorities.

1.5.8 Release Cycle

The CA:FeX Interoperability Specifications' release cycle will include a multi-stage review and feedback process, as documented [here](#).

1.6 CA:FeX Use Case Overview

1.6.1 Use Case Overview

This section describes the two Use Cases for the FHIR Health Information Exchange (HIE) Implementation Pattern, including all design constraints and assumptions as well as the flows of information that will be specified in the CA:FeX Interoperability Specifications. This section also introduces the scenarios that describe how the specified workflows may be used in the Canadian e-health context.

1.6.2 In-Scope

The following Use Cases are in scope for Release 1 of the CA:FeX Interoperability Specifications v1 Trial Implementation:

- UC-01 Create and Submit Document
- UC-02 Query and Retrieve Document

1.6.3 Use Case Actors and Services

The Use Case Actors and the Services that are used by this specification are described below. Additional information can be found in the [Core Interoperability Specifications Requirements](#) section.

Use Case Actors and Descriptions

| Actor Name | Description / Definition |
|---|--|
| Producer | A health records system (e.g., EMR, HIS, CIS, PHR, or EHR) that creates/produces clinical data in response to a request from an authorized health care provider, the subject of care or another authorized health records system. |
| Consumer | A health records system (e.g., EMR, HIS, CIS, PHR, Patient Portal or EHR) that enables access to or receipt of clinical data by an authorized health care provider or the subject of care/patient. |
| Clinical Data Repository (Local or Central) | A Clinical Data Repository is a shared storage space for clinical documents that can be hosted locally (e.g., at the data producer) or at the Central Infrastructure and can be accessed by authorized users. |
| Central Infrastructure | A Central Infrastructure collects health information from participating organizations and stores the information in a centralized place. The Infrastructure also provides access control. Typically, the Central Infrastructure is under jurisdictional control. |

Use Case Actor Mapping

| Actor Name | UC-01 | UC-02 |
|---|-------|-------|
| Producer | x | |
| Consumer | | x |
| Clinical Data Repository (Local or Central) | x | x |
| Central Infrastructure | x | x |

1.6.4 Design Constraints & Assumptions

The following design constraints and assumptions exist for the CA:FeX Interoperability Specifications:

- All Use Case Actors/Users are logged in to the system. The Use Case Actors/Users are authenticated and appropriately authorized for all data exchange transactions.
 - Our recommendation is that the Use Case Actors/Users obtain a valid access token from the Authorization Server that is used within each transaction and is based on the IUA IHE Profile.
 - Vendors or jurisdictions may decide not to implement optional IHE profiles such as IUA, ATNA and CT, however it is highly recommended that areas pertaining to authentication, auditing and security are being addressed using solutions that currently exist in their respective enterprise architecture.

1.7 Core Interoperability Specification Requirements

1.7.1 Actor Mapping to Interoperability Specification

The Use Case Actors and the Services they support are described at a functional level in the [Use Cases and Definitions](#) section of the CA:FeX Interoperability Specifications. Services may be Required or Optional. The Use Case Actor, Service(s) and optionality are conveyed in the first three columns of Tables 1 and 2 in the section below. The second part of the table (columns 4-7) provides the mapping for the Use Case Actor to the detailed specifications (such as Technical Actors and Optionality) that systems shall implement to exchange healthcare information in the context of these use cases. The purpose of the tables below is to map the use case actors to the technical actors and the services they are supporting. These tables are aligned with the flow captured in the sequence diagrams which are included in the Reference Architecture section.

For a selected Use Case Actor (a single row in the table), the system shall implement all of the requirements (some optionality when allowed) listed in the second part of the table (columns 4-7). This includes the referenced healthcare profiles, the standards specified and terminology standards. For Technical Actors, which map to CA:FeX or an IHE Profile (*CA:FeX, PDQm, PMIR, etc.*), the last column provides the reference location of the specification. Links for these referenced specifications have been included in the tables below. Additionally, the below table does not showcase all of the possible combinations of IHE profiles and transactions that can be used for a particular implementation pattern. For example, 'Identify Patient' service can use PIXm in place of PDQm if the preferred implementation pattern is PIXm/PMIR.

Versioning

This is an evolving specification; the release cycle assumes some degree of change will happen across versions. We've established a versioning protocol and will be clearly communicating the version of the artefacts that are employed by the CA:FeX Interoperability Specifications and will ensure testing aligns to it. There will be a process in place to monitor changes in the current versions of the IHE profiles which will be incorporated in future versions of the CA:FeX Interoperability Specifications based on our interoperability roadmap.

Published Versions

The following describes the published versions in scope for the required and optional IHE Profiles that have been referenced in this specification:

- [IUA](#): Revision 2.1 - Trial Implementation
- [PDQm](#): v2.3.0: Trial Implementation based on FHIR R4
- [PMIR](#): Revision 1.3 – Trial Implementation

Legend

R = Required

O = Optional

1.7.2 Table 1 Interoperability Conformance Requirements for Use Case 1: Create and Submit Document

| USE CASE 1: Create and Submit Document | | | MAPPING TO SECTIONS FROM THIS AND REFERENCED INTEROPERABILITY SPECIFICATIONS | | | |
|--|--|-----|--|-----|--|---|
| USE CASE ACTOR | SERVICE SUPPORTED | OPT | TECHNICAL ACTOR | OPT | PROFILE/ STANDARD | REFERENCED SPECIFICATION AND STANDARDS |
| Producer | Authenticate User | O | Client (e.g. EMR) | O | Internet User Assertion (IUA) | Refer to IUA : Revision 2.1 - Trial Implementation |
| | Identify Patient | O | Client (e.g. EMR) | O | Use Existing Standards Employed by the Clinical System | N/A |
| | | O | Patient Demographic Consumer | O | PDQm | Refer to PDQm : v2.3.0: Trial Implementation based on FHIR R4 |
| | Retrieve Clinical Data (Patient Identifier) | R | Client (e.g. EMR) | R | Use Existing Standards Employed by the Clinical System | N/A |
| | Assemble and Review Document | R | Client (e.g. EMR) | R | Use Existing Standards Employed by the Clinical System | N/A |
| | Omit or Mask Data based on Jurisdictional Policy | O | Client (e.g. EMR) | O | Jurisdictional Requirement | N/A |
| | Save Document to Clinical Data Repository | R | Client (e.g. EMR) | R | Use Existing Standards Employed by the Clinical System | N/A |

| USE CASE 1: Create and Submit Document | | | MAPPING TO SECTIONS FROM THIS AND REFERENCED INTEROPERABILITY SPECIFICATIONS | | | |
|--|---|-----|--|-----|-------------------|---|
| USE CASE ACTOR | SERVICE SUPPORTED | OPT | TECHNICAL ACTOR | OPT | PROFILE/ STANDARD | REFERENCED SPECIFICATION AND STANDARDS |
| | | R | Data Source | R | CA:FeX | Refer to Submit a Document in the Exchanging Documents in FHIR section. |
| Clinical Data Repository (Central) | Save Bundle to Clinical Data Repository | R | Data Recipient | R | CA:FeX | Refer to Submit a Document in the Exchanging Documents FHIR section. |
| Central Infrastructure | Identify Patient | O | Patient Identity Registry | O | PMIR | Refer to PMIR : Revision 1.3 – Trial Implementation |

1.7.3 Table 2 Interoperability Conformance Requirements for Use Case 2: Query and Retrieve Document

| USE CASE 2: Query and Retrieve Document | | | MAPPING TO SECTIONS FROM THIS AND REFERENCED INTEROPERABILITY SPECIFICATIONS | | | |
|---|-------------------|-----|--|-----|--|--|
| USE CASE ACTOR | SERVICE SUPPORTED | OPT | TECHNICAL ACTOR | OPT | PROFILE/ STANDARD | REFERENCED SPECIFICATION AND STANDARDS |
| Consumer | Authenticate User | O | Client (e.g. EMR) | O | Internet User Assertion (IUA) | Refer to IUA : Revision 2.1 - Trial Implementation. |
| | Identify Patient | O | Client (e.g. EMR) | O | Use Existing Standards Employed by the Clinical System | N/A |
| | | O | Patient Demographic Consumer | O | PDQm | Refer to PDQm : v2.3.0: Trial Implementation based on FHIR R4. |

| USE CASE 2: Query and Retrieve Document | | | MAPPING TO SECTIONS FROM THIS AND REFERENCED INTEROPERABILITY SPECIFICATIONS | | | |
|---|--|-----|--|-----|--|---|
| USE CASE ACTOR | SERVICE SUPPORTED | OPT | TECHNICAL ACTOR | OPT | PROFILE/ STANDARD | REFERENCED SPECIFICATION AND STANDARDS |
| | Request Search Document (Patient Identifier) | R | Client (e.g. EMR) | R | Use Existing Standards Employed by the Clinical System | N/A |
| | Request Document (Resource ID) | R | Client (e.g. EMR) | R | Use Existing Standards Employed by the Clinical System | N/A |
| | Return <i>searchset</i> Bundle | R | Data Consumer | R | CA:FeX | Refer to Search for a Document in the Exchanging Documents in FHIR section. |
| | Return <i>document</i> Bundle | R | Data Consumer | R | CA:FeX | Refer to Retrieve a Document in the Exchanging Documents in FHIR section. |
| Clinical Data Repository (Central) | Retrieve Resources from Clinical Data Repository | R | Data Responder | R | CA:FeX | Refer to Search for a Document in the Exchanging Documents in FHIR section. |
| | Retrieve Bundle from Clinical Data Repository | R | Data Responder | R | CA:FeX | Refer to Retrieve a Document in the Exchanging Documents in FHIR section. |
| Central Infrastructure | Identify Patient | O | Patient Identity Registry | O | PMIR | Refer to PMIR : Revision 1.3 – Trial Implementation. |

1.8 Actor Conformance

A system conforming to the CA:FeX Interoperability Specifications shall claim conformance at the level of a Use Case Actor (first columns of [Table 1](#) and [Table 2](#) of section [Core Interoperability Specification Requirements](#)). A system may claim conformance to one or more Use Case Actors among:

- Producer
- Consumer
- Clinical Data Repository (Local or Central)
- Central Infrastructure

Producer and Consumer use case actor roles will primarily be taken up by EMR clinical solution vendors. Clinical Data Repository and Central Infrastructure use case actor roles can be taken up either by EMR clinical solution vendors or jurisdictions depending on the implementation approach that the jurisdiction decides to adopt. Similarly, the Consumer can be a Patient Portal, in which case the use case actor can be taken up either by a vendor or jurisdiction depending on the approach and policies defined regarding patient / subject of care access to their personal health information (PHI).

In order to implement a system that fully supports the CA:FeX Interoperability Specifications v1 Trial Implementation, the system shall be able to demonstrate that it conforms to every required actor and transaction for which it is claiming conformance.

1.8.1 Constraints on Use Case Actors

The section below captures some of the design constraints on use case actors when developing functionality to support the services mapped to them.

Note: The scope of this section is limited to the constraints that are applicable to actors and transactions defined for CA:FeX Interoperability Specifications (See section [CA:FeX Actors and Transactions](#)). The two key services supported by CA:FeX are:

- Save Document to Clinical Data Repository
- Retrieve Document from Clinical Data Repository

The following section provides key design constraints for implementation of these two required services using RESTful APIs based on CA:FeX and FHIR standards. To support these two services, the following three RESTful transactions have been defined:

| Service Supported | FHIR RESTful Transactions |
|---|-----------------------------|
| Save Document to Clinical Data Repository | 1. Submit Data [CA:FeX-1] |
| Retrieve Document from Clinical Data Repository | 2. Search Data [CA:FeX-2] |
| | 3. Retrieve Data [CA:FeX-3] |

Save Document to Clinical Data Repository

The Producer and Clinical Data Repository Use Case Actors are required to implement the *Save Document to Clinical Data Repository* service.

These actors shall use the FHIR transaction **Submit Data [CA:FeX-1]** of CA:FeX that executes a *Submit Data request* from a Data Source to a Data Recipient.

Submit Data [CA:FeX-1]

This message involves a request by a Data Source to transfer a FHIR document to a Data Recipient. The request is received by a Data Recipient which stores the received FHIR documents and returns an HTTP response code.

Trigger Events

This method is invoked when the Data Source needs to submit one or more FHIR documents to a Data Recipient (Clinical Data Repository).

Message Semantics

This message uses the **HTTP POST** method on the target Submit Data endpoint to convey the metadata and the document(s) as a FHIR transaction. The Data Source shall initiate a FHIR “transaction” using a “create” action by sending an **HTTP POST** request method composed of a FHIR Resource. The media type of the HTTP body shall support application/fhir+json and should support application/fhir+xml. Additional information on this FHIR implementation pattern is provided in the [Submit a Document](#) section under Exchanging Documents in FHIR.

Expected Actions

The Data Recipient shall accept media types application/fhir+json and application/fhir+xml. On receipt of the submission, the Data Recipient shall validate the resources and respond with one of the HTTP response codes and an [OperationOutcome](#), if applicable. Additional information on handling responses is provided in [Response Handling](#) section under Exchanging Documents in FHIR.

Retrieve Document from Clinical Data Repository

The Consumer and Clinical Data Repository (Central) Use Case Actors are required to implement the *Retrieve Document from Clinical Data Repository* service.

These actors shall use the following transactions to find metadata and retrieval of identified clinical document:

- **Search Data [CA:FeX-2]**
- **Retrieve Data [CA:FeX-3]**

Search Data [CA:FeX-2]

This message involves a query request by a Data Consumer to find FHIR documents using parameterized queries. The request is received by a Data Recipient which returns a Bundle containing the matching search parameters.

The Data Consumer may use HTTP GET or HTTP POST based searches. The Data Responder shall support both GET and POST based searches.

Trigger Events

When a Data Consumer needs to discover a list of FHIR Bundle Resources or retrieval of documents that are anticipated to be a mix of FHIR-assembled (Composition) and FHIR-enabled (Binary) documents.

Message Semantics

The Data Consumer executes an **HTTP search request** against the Data Responder endpoint (FHIR Repository).

The Data Consumer may use **HTTP GET** or **HTTP POST** based searches. The Data Responder shall support both GET and POST based searches.

GET [base]/[resourcetype]?name=value&...

POST [base]/[type]/_search?{[parameters]}{&_format=[mime-type]}

Query Search Parameters

The following search parameters are used generically to support document retrieval.

Additional search parameters that are specific to the requirements for a specific use case or IGuide (e.g., PS-CA) may be further defined by those implementations.

| Query Search Parameters | Applied to | Description |
|-------------------------|---------------------------------------|---|
| timestamp | bundle.timestamp | This parameter, of type date, specifies the timestamp when the FHIR bundle was created. See FHIR http://hl7.org/fhir/R4/search.html#date for use of the date search type. |
| patient.identifier | bundle.composition.patient.identifier | This parameter, of type token, specifies an identifier associated with the patient to which the document is assigned. This use of patient.identifier follows the FHIR Chaining Parameters search methodology. |
| type | bundle.composition.type | This parameter, of type token, specifies the kind of document (LOINC if possible). The use of bundle.composition.type follows the FHIR Chaining Parameters search methodology. |
| status | bundle.composition.status | This parameter, of type token, specifies the status of the composition. The use of bundle.composition.status follows the FHIR Chaining Parameters search methodology. |
| author | bundle.composition.author | This parameter, of type reference, specifies who and/or what authored the document. The use of bundle.composition.author follows the FHIR Chaining Parameters search methodology. |
| date | bundle.composition.date | This parameter, of type date, specifies when this document reference was created. The use of bundle.composition.date follows the FHIR Chaining Parameters search methodology. |

Additional information on this FHIR implementation pattern including Document Search Pattern options is provided in the [Search for a Document](#) section under Exchanging Documents in FHIR.

Expected Actions

The Data Responder shall process the query and return a search result matching the search criteria included in the request. The FHIR standard provides encodings for responses as either XML or JSON.

Additional information on handling responses is provided in the [Response Handling](#) section under Exchanging Documents in FHIR.

Security Considerations

Data Responder SHALL reject any unauthorized requests by returning an HTTP 401 unauthorized response code. This transaction should not return information that the Data Consumer is not authorized to access. Where authorization here is inclusive of system, app, and user according to local policy, patient consents, and security layering. However, the transaction may return search bundle that have Reference elements that the Data Consumer may not have access to. This is to say that the authorization need only be to the content returned in the Bundle. There may be references (URLs) for which the content is not authorized. This is considered proper as the Data Consumer would need to retrieve the content pointed to by those references, and at that time the proper authorization decision would be made on that context and content. In this way it is possible for a Data Consumer to get Resources that are pointing at data that the Data Consumer is not authorized to retrieve. Thus, the URLs used must be carefully crafted so as to not expose sensitive data in the URL value.

Retrieve Data [CA:FeX-3]

This message involves a request by a Data Consumer for retrieving known documents using the document's resource id from a Clinical Data Repository. The request is received by a Data Responder which returns the requested FHIR document(s) or returns an HTTP response code.

This message uses the HTTP GET request to retrieve the identified FHIR document(s) from the central Clinical Data Repository.

Trigger Events

This method is invoked when the Data Consumer needs to retrieve FHIR document(s).

Message Semantics

The Data Consumer sends an HTTP GET request to the server based on a known resource id from the Data Responder. This operation will return a Bundle that was previously identified in a search/or prior retrieval or an embedded attachment in DocumentReference or a Binary resource.

GET [base]/[resourcetype]?name=value&...

Additional information on this FHIR implementation pattern including Document Search Pattern options is provided in the [Retrieve a Document](#) section under Exchanging Documents in FHIR.

Expected Actions

The Data Responder shall process the GET request and respond with FHIR Document (s) matching the specified identifier included in the request. When the requested document is returned, the Data Responder shall respond with an HTTP Status Code 200. The HTTP message-body shall be the content of the requested document.

Additional information on handling responses is provided in the [Response Handling](#) section under Exchanging Documents in FHIR.

Security Considerations

Data Responder SHALL reject any unauthorized requests by returning an HTTP 401 unauthorized response code. This transaction should not return information that the Data Consumer is not authorized to access.

1.9 Privacy & Security Guidance

1.9.1 Privacy Considerations

- i** Infoway is developing a privacy toolkit to be released in April, Interoperability Challenge Accepted: A Privacy Toolkit for Canadian Jurisdictions, that will address privacy considerations at a pan-Canadian level.


1.9.2 Security Considerations

- i** Fast Healthcare Interoperability Resources (FHIR) is not a security protocol, nor does it define any security related functionality. However, FHIR does define exchange protocols and content models that need to be used with various security protocols defined elsewhere. FHIR transactions defined as part of the CA:FeX implementation pattern often make use of patient-specific information which could be exploited by malicious actors resulting in exposure of patient data. For this reason, all FHIR transactions must be secured appropriately with access to limited authorized individuals, data protected in transit, and appropriate audit measures taken. Implementers **SHOULD** be aware of security considerations associated with FHIR transactions (<http://hl7.org/fhir/R4/security.html>), particularly those related to:
- Communications
 - Authentication
 - Authorization/Access Control
 - Audit Logging
 - Digital Signatures
 - Security Labels
 - Narrative

Additionally, many FHIR transactions using HTTP REST will include query parameters that would be identifiers, quasi-identifiers, or sensitive health topics. For example, it is common for patient identifier to be a query parameter. With this URL pattern, the query parameters are typically visible in the server audit log or browser history. The risk from this visibility should be mitigated in system or operational design, by protecting the logs as sensitive data, or by designing other measures into the system to prevent inappropriate exposure.

1.10 CA:FeX Actors and Transactions

The following diagram illustrates the actors and transactions in scope for the CA:FeX Interoperability Specification.



There are three primary transactions for the CA:FeX Interoperability Specification, along with additional sub-transactions listed below for specific circumstances. Please refer to the [Exchanging Documents in FHIR](#) section for more details.

| Transaction ID | Description |
|----------------|---|
| CA:FeX-1 | Submit Data |
| CA:FeX-2 | CA:FeX-2A: Search Against FHIR Assembled Documents Repository |
| | CA:FeX-2B: Search Against Hybrid Documents Repository |
| CA:FeX-3 | CA:FeX-3A: Retrieve Document From FHIR Assembled Documents Repository |
| | CA:FeX-3B: Retrieve Document from Hybrid Documents Repository |

2 Use Cases and Definitions

The purpose is to describe the use cases and workflow scenarios for sharing clinical documents across solutions. Each jurisdiction may have implementation variances within the use cases. Therefore, these use cases do not represent required implementation choices, nor are they representative of all possible implementation choices. These use cases are to be considered as examples. The use cases provide high-level interactions between a Health Care Provider or Patient using a Clinical Solution (e.g., EMR, Patient Portal, etc.) and the HIE. Use cases provide the business description or "conversation" between the system(s) and its user(s), known as Participants. Participants can be people (e.g., health care providers, patients, etc.) using a Clinical Solution (e.g., EMR, Patient Portal) or systems (e.g., HIE, etc.). Please note that detailed interactions are defined in the Reference Architecture section of this document.

Each use case will include:

- use case scenario,
- examples of use case triggers, pre- and post-conditions,
- who the participants are (i.e., people and systems),
- a use case diagram to provide a visual representation of the interactions between participants,
- use case steps corresponding to the diagram and potential alternate flows, and
- requirements.

2.1 Use Case Index

This section includes a proposed list of use cases which were identified as being priority use cases in the pan-Canadian environmental scan. Additional use cases will be defined in future releases.

The scope for this release of the CA:FeX Specifications has been defined to include use cases UC-01 and UC-02.

The list below includes the use cases a Use Case ID, name and description of the use case.

| Use Case ID | Use Case Name | Use Case Description |
|-------------|-----------------------------|--|
| UC-01 | Create and Submit Document | A Health Care Provider, in any care setting, adds Patient clinical information for use at point of care, which is made available to other authorized HCPs. |
| UC-02 | Query and Retrieve Document | Query and retrieval of clinical documentation performed by a Health Care Provider for use at the point of care or by the Patient themselves to obtain a copy of their own personal health information. |

2.2 Requirement Priority Definitions

| Priority | Definition |
|----------|---|
| SHALL | <ul style="list-style-type: none"> • used to indicate a required requirement. |

| Priority | Definition |
|----------|--|
| SHOULD | <ul style="list-style-type: none">used to indicate that a requirement is recommended and should be considered as best practice for implementation, but not required (i.e., it is optional) for implementation. |

2.3 UC-01 Create and Submit Document

Description

A Health Care Provider, in any care setting, adds Patient clinical information for use at point of care, which is made available to other authorized HCPs.

Scenario

A patient schedules a visit with their regular health care provider, within their Medical Home, with symptoms including dizziness and an earache. The patient mentions that since they last visited, another clinic noted that they have high blood pressure (hypertension) which is being monitored at home for now. The patient also mentions a suspected penicillin allergy. The health care provider determines that the patient has an external ear infection (otitis externa) and prescribes antibiotics. The health care provider creates a clinical note in their EMR, which may trigger automatic updates, such as updates to the prescription information. The health care provider decides to submit this new information to the network (i.e., Clinical Data Repository) so that it is available for other health care providers who may be providing care for this patient.

Triggers, Pre-conditions, Post-Conditions

This section describes example triggers, pre-conditions & post-conditions related to uploading new clinical information to the Clinical Data Repository. It is not inclusive of all potential workflow scenarios which may be implemented within Canadian jurisdictions.

Triggers

- Health Care Provider provides care to a patient and adds clinical information to the Patient's record.
- Health Care Provider receives additional information for a patient that they wish to share with other HCPs. For example, HCP receives test results for a Patient and adds the clinical information to the Patient's problem list.

Pre-conditions

- Clinical documentation shall uniquely identify the Patient so that it can be uploaded to the Clinical Data Repository and available to other HCPs (e.g., uniquely identified by a Client Registry ID)
- In jurisdictions where explicit consent is required to share Patient clinical information:
 - Patient provides, or has previously provided, consent to share their data to the Clinical Data Repository.

Post-conditions

- New Patient clinical information recorded/registered in the Clinical Data Repository.
- Authorized health care providers have access to view the Patient's clinical information or may receive a notification that new clinical information about the patient is available.

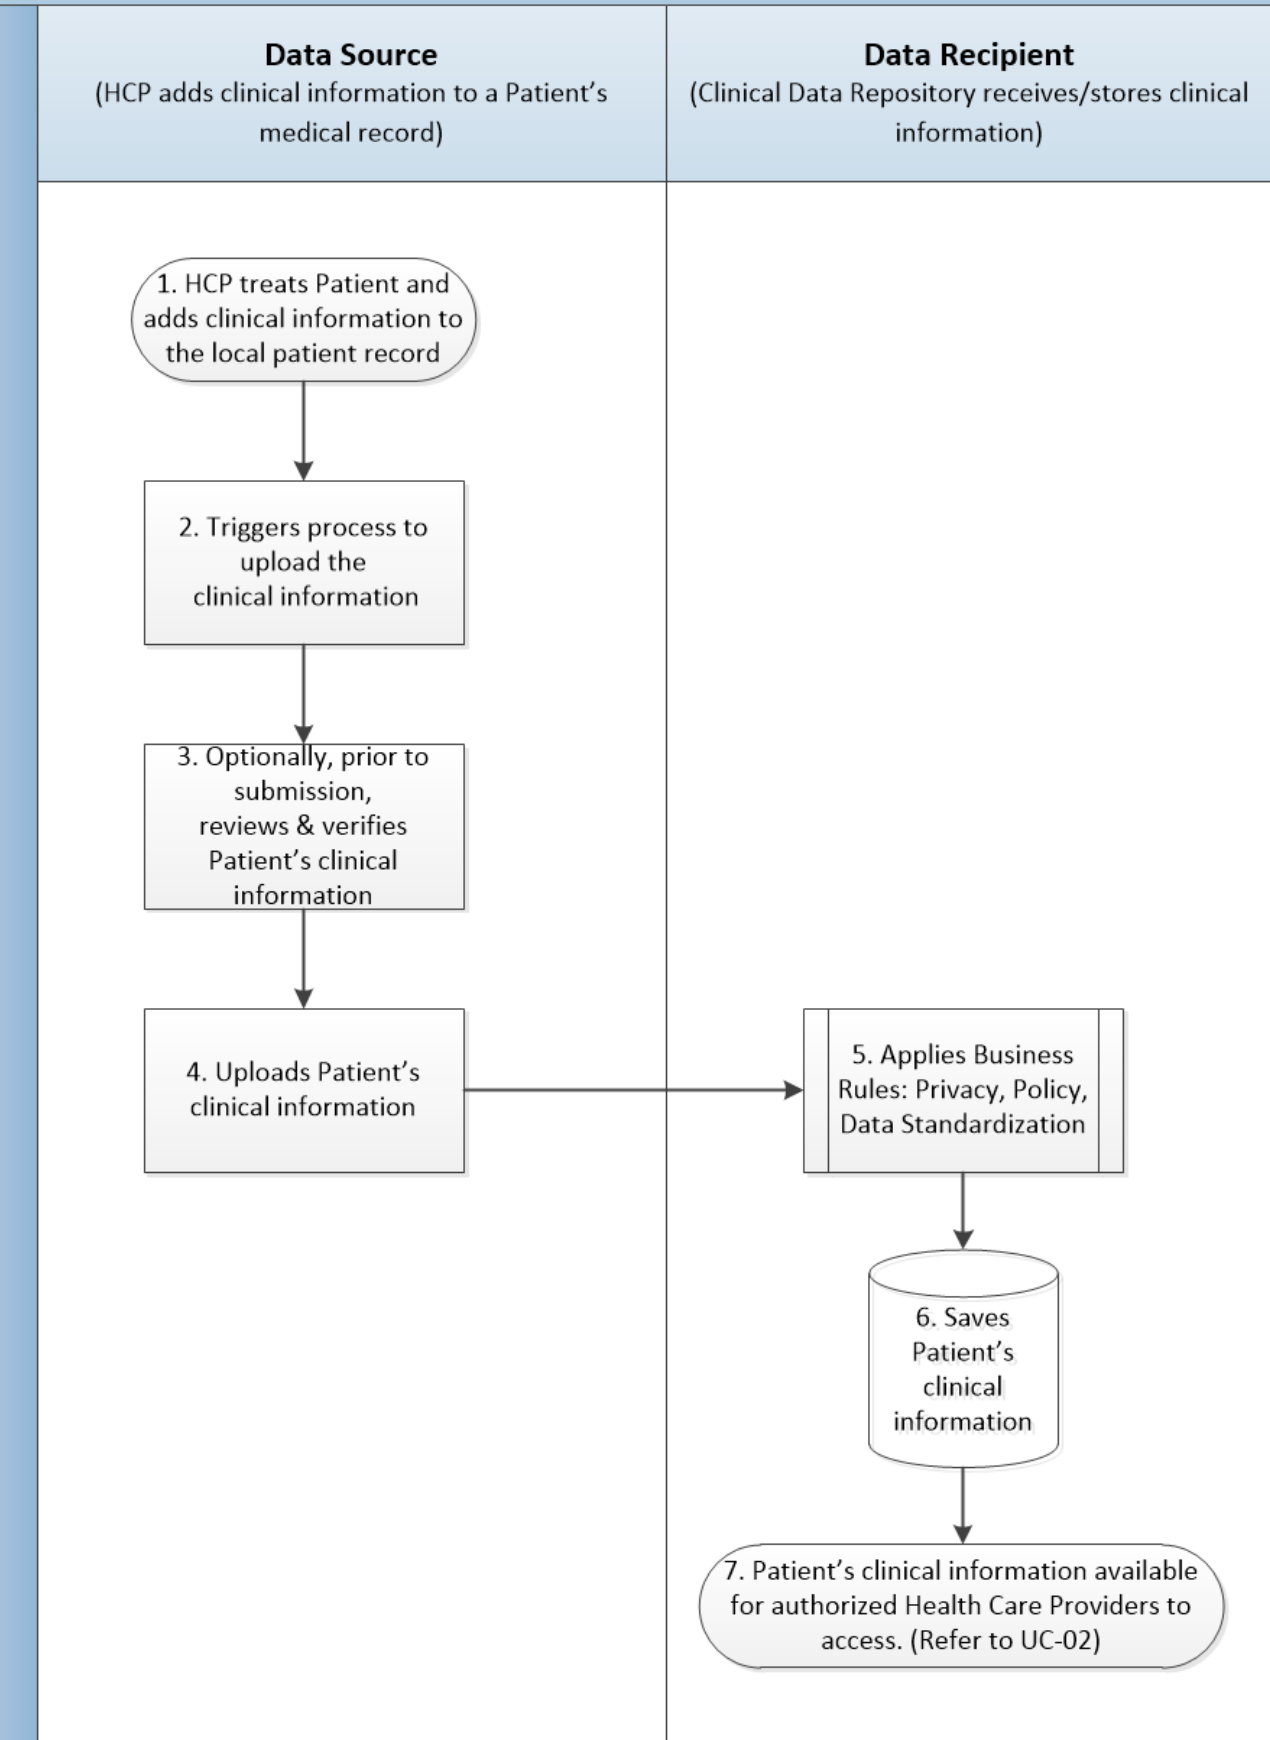
Use Case Participants & Diagram

The participants involved in this use case are:

- Data Source (Health Care Provider adding patient clinical information via their local Clinical Solution (e.g., EMR))
- Data Recipient (Clinical Data Repository receiving/storing the Patient's clinical information)

This use case diagram represents the participants and their role in the use case with a high-level view of the flow of information.

UC-01 Health Care Provider Uploads Health Information



Use Case - Primary Flow

The following provides a textual description corresponding to the use case diagram.

1. Health Care Provider treats Patient and adds clinical information to the Patient's health record in their local Clinical Solution (e.g., EMR, HIS).
2. Health Care Provider decides to share the new clinical information collected and triggers the process to upload the information to the Clinical Data Repository.
3. Health Care Provider, optionally, reviews and validates the Patient's clinical information prior to sharing/ uploading it to the Clinical Data Repository.
4. Health Care Provider sends / uploads the Patient's clinical information to the Clinical Data Repository.
5. Clinical Data Repository applies business rules (e.g. data standardization, privacy, policy, etc.).
For example:
 - a. Validation of Patient Summary data (e.g. Provider identified and eligible to submit clinical information, Patient identified, etc.)
 - b. Checks for existing clinical information for same patient/same provider - apply replacement / archiving rules
6. Clinical Data Repository saves the Patient's clinical information.
 - a. Clinical Data Repository responds to the submitting system to indicate that the submission was accepted (i.e., recorded/registered) or it was a bad request.
7. Patient's clinical information is available for access by authorized Health Care Providers. (Refer to UC-02 Query and Retrieve Document)

Use Case - Alternate Flow

The following list provides possible alternate flows that may occur within this use case.

- Step 3: Health Care Provider has the option to bypass an additional review of the clinical information, allowing the Clinical Solution to automatically share/upload the Clinical Information to the Clinical Data Repository.
- Step 3: Health Care Provider, upon review of the clinical information, chooses to make changes within the local Patient's health record prior to uploading it to the Clinical Data Repository.
- Step 4: Health Care Provider, after submitting the Patient's clinical information, identifies that there is incorrect or missing information. The HCP will have the option to modify and upload the corrected clinical information to the Clinical Data Repository.
- Step 4: Health Care Provider, after submitting the Patient's clinical information, identifies that incorrect information has been uploaded (e.g., information is for the wrong patient). The HCP will have the option to retract / delete the information from the Clinical Data Repository.

Use Case - Requirements

The following is a list of key requirements that will be addressed as part of this use case.

| # | Category | Requirement Description |
|---|----------|---|
| 1 | Write | FHIR API SHALL be capable of accepting write operations to allow creation of new clinical document in the central Clinical Data Repository |
| 2 | Response | FHIR API SHALL be capable of returning a response that a new clinical information has been successfully recorded/registered in the central Clinical Data Repository |

| # | Category | Requirement Description |
|---|----------|---|
| 3 | Response | FHIR API SHALL be capable of returning a response that a Patient clinical information could not be recorded/registered in the central Clinical Data Repository (HTTP 400 - Bad Request) |

2.4 UC-02 Query and Retrieve Document

Description

Query and retrieval of clinical documentation performed by a Health Care Provider for use at the point of care or by the Patient themselves to obtain a copy of their own personal health information.

Scenario

The following are example scenarios. It is not inclusive of all potential scenarios which may be implemented within Canadian jurisdictions.

1) A Health Care Provider, in any care setting, queries and retrieves a clinical document for use at the point of care

A patient schedules a visit with a health care provider, outside of their Medical Home, with symptoms including dizziness and an earache. The patient mentions that they have a regular health care provider, within their Medical Home, and experiences high blood pressure (hypertension) which is being monitored at home for now. The health care provider collects information from the patient and, using their Clinical Solution (e.g., EMR), searches for clinical information (e.g., searches the network to locate clinical documentation created and shared by another Health Care Provider). Upon finding a clinical document(s) for their patient, the health care provider views and uses the information in support of providing care for this patient.

2) A Patient or Subject of Care accesses/views and can obtain a copy of their own personal health information.

A patient, or their designated caregiver, would like to access their personal health information (PHI) to stay up to date with their medical health information, empowering them to play an active role in their own care.

Triggers, Pre-conditions, Post-conditions

This section describes example triggers, pre-conditions & post-conditions related to the query and retrieval of clinical document(s) from the Clinical Data Repository. It is not inclusive of all potential workflow scenarios which may be implemented within Canadian jurisdictions.

Triggers

Scenario 1:

- Patient visits Health Care Provider for care.
- Where applicable, HCP received a notification that new clinical information is available for the Patient to which they have subscribed to receive notifications.

Scenario 2:

- Patient, or their designated caregiver, chooses to view personal health information to stay informed of their medical information.
- Patient wants to obtain a copy of their personal health information to have on their person while travelling.
- Patient wants to obtain a copy of their personal health information to share with another care provider.

Pre-conditions

Scenario 1:

- Health Care Provider is logged in to their Clinical Solution (e.g., EMR).

- Health Care Provider's Clinical Solution is connected / part of the Clinical Data Repository network.

Scenario 2:

- In jurisdictions where a patient may have applied consent directives to their clinical information, HCP complies with local/jurisdictional privacy policies.
- A jurisdictional clinical system with patient access is available.
- If applicable, patient has designated and authorized a designated caregiver to access their personal health record on their behalf.

Post-conditions

Scenario 1:

- Health Care Provider views and uses the clinical document(s) in support of Patient care.

Scenario 2:

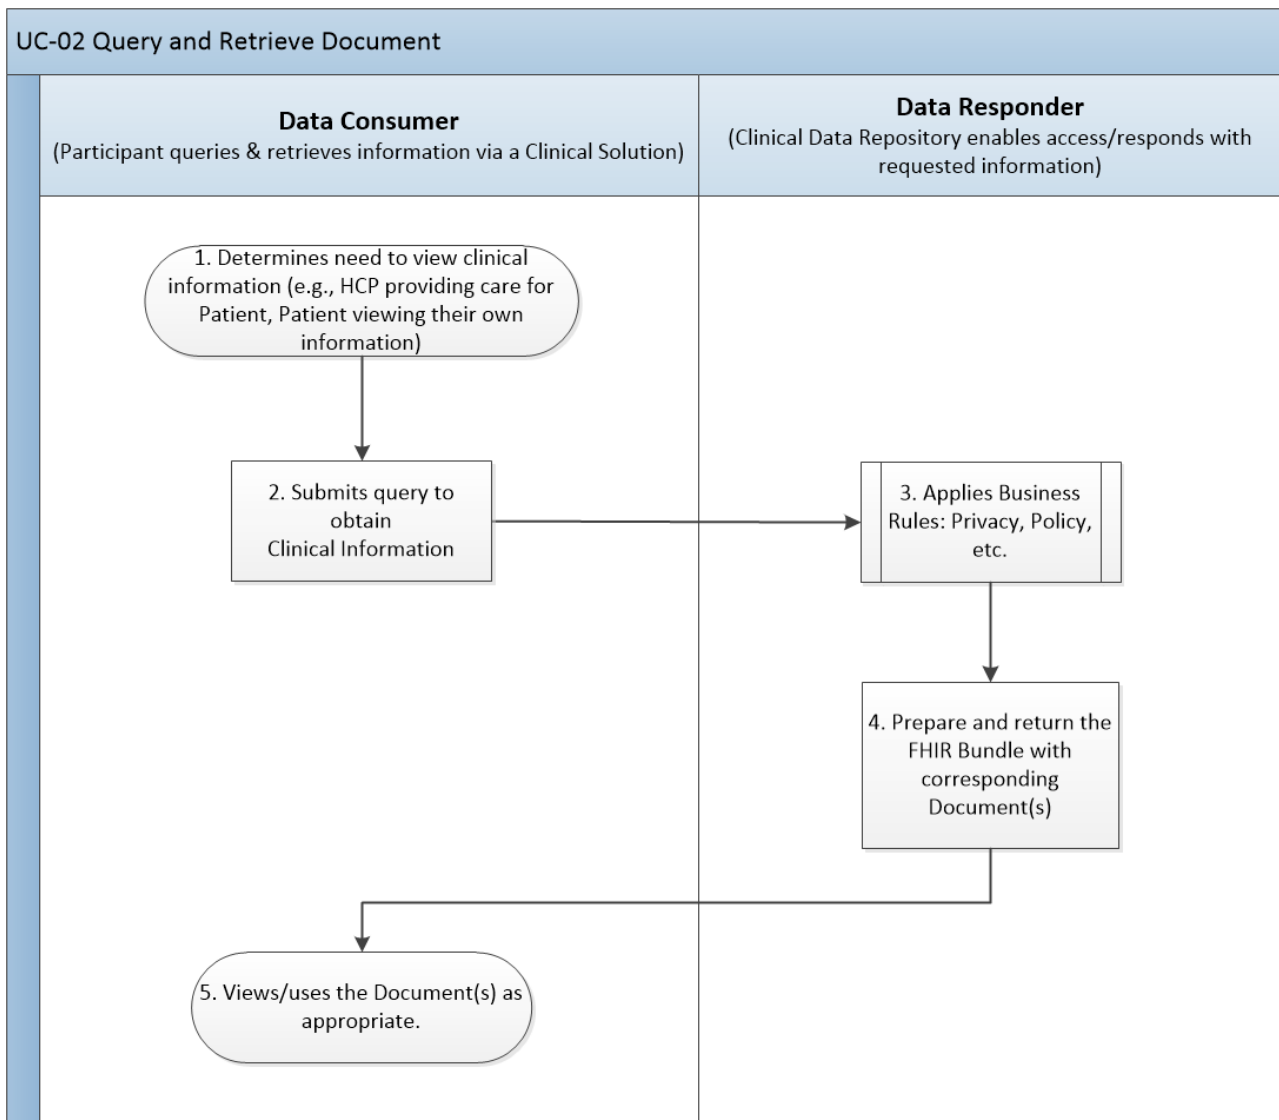
- Patient, or their designated caregiver, accessed/viewed, and optionally printed a copy of, their personal health information.
- Patient, or their designated caregiver, presents their personal health information to another health care provider to support continuity of care.

Use Case Participants & Diagram

The participants involved in this use case are:

- Data Consumer (Clinical Solution, e.g., EMR used by the Health Care Provider to request & retrieve access to clinical document(s) from the Clinical Data Repository; Patient Portal used by the Patient / Subject of Care to request & retrieve access to their clinical document(s) from the Clinical Data Repository)
- Data Responder (Health Records System acting as the Clinical Data Repository enabling access and responding with the requested clinical document(s))

This use case diagram represents the participants and their role in the use case with a high-level view of the flow of information.



Use Case - Primary Flow

The following provides a textual description corresponding to the use case diagram.

1. Participant (i.e., Health Care Provider or Patient / Subject of Care) determines need to view clinical information.
2. Participant requests access (i.e., queries) for clinical information via the clinical solution.
3. Clinical Data Repository applies applicable business/policy rules (e.g. validates requestors credentials).
4. Clinical Data Repository prepares and returns the FHIR bundle with corresponding clinical document(s).
5. Participant views/uses the clinical document(s) in support of care for the Patient or themselves.

Use Case - Alternate Flow

The following list provides possible alternate flows that may occur within this use case.

- Step 3: Patient Consent Services are not currently supported in the CA:FeX specifications. However, Patient consent is considered as a roadmap item and would be represented in a separate use case at that time where a Patient has identified consent directives requiring the Health Care Provider to address prior to

accessing the Patient's clinical information. It is recognized that some jurisdictions may have existing Patient Consent Services implemented and should be considered when implementing with that jurisdiction.

Use Case - Requirements

The following is a list of key requirements that will be addressed as part of this use case.

| # | Category | Requirement Description |
|---|----------|---|
| 1 | Query | FHIR API SHALL be capable of executing searches based on a Patient Identifier |
| 2 | Query | FHIR API SHALL be capable of executing searches based on a specified date range |
| 3 | Query | FHIR API SHOULD be capable of executing searches based on specific Document Type that the health care provider is interested in retrieving |
| 4 | Query | FHIR API SHALL be capable of accepting requests originating from a patient portal or any other system that allows a patient to initiate a request to access their personal health information (PHI) |
| 5 | Response | FHIR API SHALL be capable of responding to requests for retrieval of clinical information for requests initiated by patient to access their PHI |

3 Exchanging Documents in FHIR

i The transactions presented below represent an initial solution to the identified use cases and are expected to evolve, particularly as we hear from the implementer community and use cases are refined and added. We encourage the community to provide feedback on these emerging specifications.

3.1 Overview

HL7 characterizes a document by the following properties:

- Persistence – Documents are persistent over time. The content of the document does not change from one moment to another. A document represents information stored at a single instance in time.
- Wholeness - A document is a whole unit of information. Parts of the document may be created or edited separately, or may also be authenticated or legally authenticated, but the entire document is still to be treated as a whole unit.
- Stewardship - A document is maintained over its lifetime by a custodian, either an organization or a person entrusted with its care.
- Context - A clinical document establishes the default context for its contents.
- Potential for authentication - A clinical document is an assemblage of information that is intended to be legally authenticated.

This specification defines a document as a record (written, printed or electronic) that compiles information about a patient's care for a particular purpose or workflow. Documents are distinct from the medical data that may be included in the document as they provide an additional layer of context and metadata around how and why the pieces of medical data were compiled.

Examples of documents include discharge summaries, history & physical notes, medical imaging reports, patient summaries, and more.

Scope

This specification is starting with the exchange of documents (and their respective resources) in FHIR. As the CA:FeX Interoperability Specifications evolve, it will be expanded to tackle the exchange of individual FHIR resources that can be exchanged independent of a document workflow (e.g., retrieval of last 3 years of dispensed medications).

While HL7 FHIR defines a document format for documents that are authored and assembled in FHIR (See [FHIR Composition Resource](#)), this specification also acknowledges that FHIR can be used to exchange documents that are constructed in other formats (e.g., [FHIR Binary](#)).

The PS-CA (pan-Canadian Patient Summary) is being used as the first type of document to identify possible transaction patterns and requirements that could be extrapolated to other types of documents.

The CA:FeX Interoperability Specifications will provide detailed information for the following transactions.

| Transaction ID | Description |
|----------------|-------------|
| CA:FeX-1 | Submit Data |

| Transaction ID | Description |
|----------------|---|
| CA:FeX-2 | CA:FeX-2A: Search Against FHIR Assembled Documents Repository |
| | CA:FeX-2B: Search Against Hybrid Documents Repository |
| CA:FeX-3 | CA:FeX-3A: Retrieve Document From FHIR Assembled Documents Repository |
| | CA:FeX-3B: Retrieve Document from Hybrid Documents Repository |

3.2 FHIR Version

FHIR content in this specification is based on on [FHIR Release 4 \(R4\)](#).

3.3 Submit a Document

Scope

This capability is currently scoped to the ability to submit new FHIR Documents.

While there are a variety of ways documents can be modeled and exchanged in FHIR, the initial release of this specification is focused on the submission patterns used by early implementers that are expecting the PS-CA (pan-Canadian Patient Summary) to be delivered in FHIR Document form. Patient Summaries are a relatively new document type (e.g., legacy implementations not expected) and have been modeled as FHIR Documents.

FHIR Documents are Bundle resources with the type of "document" and which contain a Composition resource as the first entry, and key referenced resources as additional Bundle entries. For more on FHIR Documents, see [Documents](#).

3.3.1 Out of Scope

As this specification evolves, and new document types and architectures are included, this scope will also evolve to supply guidance for those business cases.

Global, pan-Canadian, and jurisdictional practices for document lifecycle management of patient summaries are still in development. For this reason, the management, verification, replacement and deprecation of documents, are out of scope of the guidance included in this release but have been included in the roadmap for future releases.

Note: This does not preclude or prevent early implementers from defining their document management practices and beginning to exercise them in their own specifications.

Info

Implementers expecting to support use cases involving multiple patient summaries over time are still expected to familiarize themselves with the boundaries the FHIR standard applies to document immutability. Once assembled into a document bundle, the document is immutable - its content can never be changed, and the document id can never be reused. Note that the document may be represented in either XML or JSON and interconverted between these or have its character encoding changed, all the while remaining the same document. However, the directly referenced content within the document and the presentation of the document cannot change substantially (such that it changes the clinical meaning of the content). Any additional documents derived from the same composition SHALL have a different document id. See <https://www.hl7.org/fhir/documents.html#content> for more details.

3.3.2 Use Cases

The following CA:FeX use case leverages this capability:

- UC-01 Create and Submit Document

3.3.3 CA:FeX Transactions

This capability describes the FHIR implementation of the CA:FeX Submit Data (CA:FeX-1) transaction found in the [Sequence Diagram for UC-01: Create and Submit Document](#).

3.3.4 HTTP Operations

This capability can be leveraged using HTTP POST requests. Implementers are expected to familiarize themselves with how create interactions are executed using the [FHIR RESTful API framework](#).

For more information on the supported HTTP operations, refer to the section [CA:FeX Actors and Transactions](#).

HTTP POST

A create is executed by performing a POST operation in the RESTful framework:

```
POST [base]/[type]{?_format=[mime-type]}
```

Content Types and Encodings:

Either `application/fhir+json` or `application/fhir+xml` content types are permitted

3.3.5 Document Submit Patterns

The optimal document submission pattern for any given implementation is determined by a number of factors. These factors include (but are not limited to): presence/absence of XDS and XCA architecture, organizational short-term and long-term goals for FHIR, current assets and prioritized capabilities, role in the health exchange ecosystem, readiness of contributing actors to adopt FHIR, etc.

These factors will be discussed in more detail in a FHIR Document Exchange Companion Whitepaper that will be published in a later release.

The patterns discussed below only apply to Data Recipients that are expecting Data Sources to submit documents using FHIR endpoints.

Some implementers may choose to use the data received through existing workflows and transform it into FHIR, particularly if their electronic medical record systems are not ready to exchange data in FHIR. Other implementations may allow for a FHIR Document to be created "on-demand" in response to a query - both are perfectly legitimate practices that do not require a submission endpoint. *These creation patterns will be surfaced as a separate capability in a later release.*

Direct Submissions of FHIR Document Bundle

Data Recipients that want to receive FHIR Documents from Data Sources need to at a minimum support the submission of the FHIR Document Bundle to a HTTP POST request to a `[base]/Bundle` endpoint.

This will [create](#) the new Bundle resource in a server-assigned location with a server assigned id for the Bundle resource. (*Note: this transaction pattern does not support [conditional create](#) which allows the resource to be created only if it does not already exist on the server.*)

Per the FHIR standard on [HTTP create interaction](#):

- The request body SHALL be a valid FHIR Resource (*in this case, a FHIR Release 4 Bundle resource*)
- The resource does not need to have an `id` element (this is one of the few cases where a resource exists without an `id` element). If an `id` is provided, the server SHALL ignore it.
- If the request body includes a `meta`, the server SHALL ignore the existing `versionId` and `lastUpdated` values. The server SHALL populate the `id`, `meta.versionId` and `meta.lastUpdated` with the new correct values
- Servers are allowed to review and alter the other metadata values, but SHOULD refrain from doing so (see [metadata description](#) for further information).

- A server SHOULD otherwise accept the resource as submitted when it accepts the create, and return the same content when it is subsequently read. However, some systems might not be able to do this; see the note on [transactional integrity](#) for discussion.
- The server returns a 201 Created HTTP status code, and SHALL also return a Location header which contains the new [Logical Id](#) and [Version Id](#) of the created resource version. See [Response Handling](#) for more details.


Considerations for Document Submission Using Bundle Endpoints

Two key attributes of FHIR Documents submitted to the Bundle endpoint are (with limited exceptions):

- Self-contained—All of the primary resources used in the document must be included in the document (see [Composition](#) in the FHIR standard); other resources used in the document should also be included.
- Point-in-time—The entire document is stored under the /Bundle endpoint, and the content is not (generally) updated. If the Data Receiver actor, which supports the /Patient endpoint outside of this transaction, updates an address in a Patient resource under /Patient, a Patient resource inside a Document would not be expected to reflect that.

Submission of a FHIR Document to a Bundle endpoint is not enough on its own for this pattern be effectively used for document lifecycle management. This pattern is seen as a "building block" that can be further augmented either through a) API interactions and operations, or through b) implementation of significant internal business logic to make determinations on document currency for example.

Note: Despite the fact that a FHIR Document may contain Composition, Patient, Practitioner or other resources, posting a FHIR Document Bundle to a Bundle Endpoint does not make those resources automatically available at their equivalent endpoints (e.g., /Composition, /Patient, /Practitioner, etc.).

-  Outside of the scope of this specification's guidance, an implementation may choose to decompose a received FHIR Document, and manage the included resources individually. An implementation that does this should consider several issues, including:
- Resource deduplication (e.g., when the Patient or Composition resource in a document is actually the same as an already known Patient or Composition resource, perhaps obtained from another document)
 - Sources of truth for various resources and document currency (e.g., whether to update an existing patient address based on a 10-year old document)
 - Whether and how updates to the extracted resources feed back into the submitted document
 - Document lifecycle compared to resource lifecycle (i.e., if a Practitioner resource is removed from the system 5 years after the Practitioner retires, is a FHIR Document or the Composition extracted from it still usable, etc.)
 - Whether to persist the submitted document as well as the extracted resources

3.3.6 Document Submit Patterns: Augmenting Base Capabilities

The section will evolve in future releases to cover additional submission patterns for documents (e.g., submitting a binary version of the document, performing a bundle transaction, submitting supporting resources that aid in the finding of the document, etc.).

3.3.7 Cutting Edge Capabilities

FHIR Operations

FHIR Operations are used (a) where the server needs to play an active role in formulating the content of the response, not merely return existing information, or (b) where the intended purpose is to cause side effects such as the modification of existing resources, or creation of new resources.

This specification does not currently define any custom operations for the Document Submit Capability. However, a number of operations exist that are actively being tested by the international standards community that can support the creation of a document. These have been listed below for implementers to test and provide feedback on.

Generate a Document Bundle using the \$document Operation - This operation on a Composition (at /Composition/[id]) is helpful for systems that want to generate (and optionally, persist) a document Bundle from an existing Composition, including those that do not retain an original Document used to generate a decomposed Composition.

- Compositions stored at the /Composition endpoint are not FHIR Documents, although the \$document operation may be used to create point-in-time FHIR Documents from a Composition. This operation has been surfaced in the hyperlink above but is not an official supported document creation pattern in this release. Feedback from implementers on the use of \$document is requested.

3.4 Search for a Document

3.4.1 Scope

This capability supports the ability to find documents using parameterized queries.

There are two types of query behavior this capability currently supports.

- Searches that are intended to return a payload that includes any FHIR Documents that match the parameters provided. These types of queries combine search and retrieval but put more strain on the requesting system to manage and filter the volume of content in the response in a way that is more manageable to the user.
- Searches that are intended to create a FHIR Document or a pointer to a document (e.g., DocumentReference) on demand if the repository has content that matches the parameters provided. These types of searches leverage FHIR operations and tend to reduce the burden on requesting systems to have to tailor their request for a document to include the format the document was stored in. This makes this type of search more ideal for querying hybrid repositories.

Note: Details on what the responder supports can be found within the CapabilityStatement that the requester retrieves as a separate step from this search. Querying for a CapabilityStatement is an expected pre-requisite for being able to query a source. However, it is not required to be conducted at the beginning of every session to avoid forcing performance costs on clients that interact with known servers that do not update their security services and/or supported capabilities frequently.

3.4.2 Use Cases

The following CA:FeX use cases leverage this capability:

- [UC-02 Query and Retrieve Document](#)

3.4.3 CA:FeX Transactions

This capability describes the FHIR implementation of the CA:FeX Search Data (CA:FeX-2) transaction found in the [Sequence Diagram for UC-02: Query and Retrieve Document](#).

3.4.4 HTTP Operations

This capability can be leveraged using HTTP GET requests and/or HTTP POST. Implementers are expected to familiarize themselves with how searches are executed using the [FHIR RESTful API framework](#).

See [CA:FeX Actors and Transactions](#) for more information on the supported HTTP operations.

HTTP GET

In the simplest case, a search is executed by performing a GET operation in the RESTful framework:

```
GET [base]/[resourcetype]?name=value&...
```

HTTP POST

For this RESTful search (see definition in RESTful API), the parameters are a series of name=[value] pairs encoded in the URL or as an application/x-www-form-urlencoded submission for a POST:

POST [base]/[type]/_search{?[parameters]}{&_format=[mime-type]}

3.4.5 Considerations for Document Search Patterns

While there is more space for variability of document submission patterns based on a number of factors (that will be discussed in more detail in a future FHIR Document Exchange Companion Whitepaper), variations in supported document search patterns incur a higher cost towards Data Consumers.

A primary driver of this specification is to fill the gap that "FHIR-first" implementers are experiencing with existing document exchange standards like MHD and XDS. See *pan-Canadian FHIR Exchange (CA:FeX) Interoperability Specifications: Preface*. One of these particular challenges, is the need to balance efficacious and consistent search patterns against the additional effort incurred to supply supplementary metadata/resources (e.g., DocumentReference) that include more potent attributes to aid in targeted searches.

Clinical systems in US markets are more familiar with using MHD-like patterns to exchange C-CDAs using FHIR - largely due to the [reinforcement of this pattern in US Core](#). Naturally, these systems are considering how to retrofit their existing configurations to support searches for other types of documents in a similar manner.

Systems that serve non-US markets (or new solutions built entirely in FHIR for document exchange) may not have the same historical context to immediately catalyze investment of effort to supply additional resources for more desirable search behaviors. Recently, new operations have been developed that are more lightweight and may ease the burden of supporting these more potent search attributes.

Given that these patterns are evolving, and their place in the Canadian market is yet to be evaluated, CA:FeX seeks to surface these patterns and educate implementers on the values and limitations of each - before assessing which is getting the most traction from Canadian implementers.

These patterns are differentiated below using an enumerated convention (e.g., 2A, 2B).

3.4.6 Document Search Patterns: CA:FeX-2A (Search Against FHIR Assembled Documents Repository)

This pattern is intended to represent the most straightforward method of searching for FHIR Documents from a repository that does not have an existing XDS/MHD architecture and has not further augmented their search capabilities (see CA:FeX-2B).

Implementers who are trialing this pattern will either play the role of a Data Consumer or a Data Responder in an interaction where a query is constructed (using the parameters identified below) and then submitted to a [base] / Bundle endpoint using a HTTP GET or POST command.

This will execute a search which will return the results in the HTTP response as a searchset Bundle containing any document Bundle resources that match the criteria of the search. While the endpoint is anchored around Bundle, the expectation is that searches will use chaining on a handful of parameters under bundle.composition, to narrow the search to FHIR Documents.

Per the FHIR standard on [HTTP search interaction](#):

- Because of the way that some user agents and proxies treat GET and POST requests, in addition to the get based search method servers that support *search* SHALL also support a POST based search
- Supporting GET means that PHI (Personal health information) might appear in search parameters, and therefore in HTTP logs. For this reason logs should be regarded as being as sensitive as the resources themselves. This is a general requirement irrespective of the use of GET - see the [FHIR security page](#) for further commentary.

- If the search succeeds, the server SHALL return a 200 OK HTTP status code and the return content SHALL be a [Bundle](#) with `type = searchset` containing the results of the search as a collection of zero or more resources in a defined order.
- If the search fails (cannot be executed, not that there are no matches), the return value SHALL be a status code 4xx or 5xx with an `OperationOutcome`. See [Response Handling](#).

Implementers are encouraged to review the Base FHIR Specification for further details on the [HTTP search interaction](#).

Considerations for Document Search Using Bundle Endpoints

Anchoring searches around Bundle in this search pattern is recommended for implementers that want to retrieve [all the FHIR Documents](#) that meet their search criteria.

- Implementers that use this approach should consider using a combination of search parameters to reduce the processing load on responders and requesters.
- The result collection can be long, so servers may use paging. If they do, they SHALL use the method [described here](#) for breaking the collection into pages if appropriate.

This search pattern does not separate the act of search from retrieval. Clients that want to provide their users the ability to review high level information (lists, summaries, etc.) for matching documents to support selective retrieval of the desired FHIR Document bundles will need to utilize different search patterns or manage this behavior in the client's user interface.

Note: Due to limitations in querying for discrete resources when the document is loaded as a FHIR Document bundle, performing searches against endpoints for Composition or other contained resources (e.g., using Patient resource id) will only work for repositories that a) decompose the resources within the submitted bundles, or b) enable the submission of the contents of the document at their individual endpoints or through a transaction bundle so that they might be retrievable. This approach requires significant caution - implementers should be well versed in the [Document Handling Obligations](#) before considering either approach. Guidance on both mechanisms is outside the scope of this specification.


Supported Search Parameters for Document Search Using Bundle Endpoints

The following search parameters are used generically to support document retrieval anchored around the Bundle endpoint. *These will continue to evolve as new use cases and requirements are identified.*

Additional search parameters that are specific to the requirements for a particular use case or IGuide may be further defined by those implementations. See *Option 2 - FHIR Health Information Exchange (HIE) Pattern Using CA:FeX* in the pan-Canadian Patient Summary (PS-CA) Specification for details and examples of which of the parameters below are used to retrieve patient summaries.

| Query Search Parameters | Applied to | Description |
|-------------------------|------------------|---|
| timestamp | bundle.timestamp | This parameter, of type date, specifies the timestamp when the FHIR bundle was created. See FHIR http://hl7.org/fhir/R4/search.html#date for use of the date search type. |

| Query Search Parameters | Applied to | Description |
|-------------------------|---------------------------------------|---|
| patient.identifier | bundle.composition.patient.identifier | This parameter, of type token, specifies an identifier associated with the patient to which the document is assigned. This use of patient.identifier follows the FHIR Chaining Parameters search methodology. |
| type | bundle.composition.type | This parameter, of type token, specifies the kind of document (LOINC if possible). The use of bundle.composition.type follows the FHIR Chaining Parameters search methodology. |
| status | bundle.composition.status | This parameter, of type token, specifies the status of the composition. The use of bundle.composition.status follows the FHIR Chaining Parameters search methodology. |
| author | bundle.composition.author | This parameter, of type reference, specifies who and/or what authored the document. The use of bundle.composition.author follows the FHIR Chaining Parameters search methodology. |
| date | bundle.composition.date | This parameter, of type date, specifies when this document reference was created. The use of bundle.composition.date follows the FHIR Chaining Parameters search methodology. |

 There are additional parameters that can be provided in the request that help further order and organize the resources that are returned. These types of parameters are identified in the base specification as [standard parameters](#) and [result parameters](#). CA:FeX does not currently require the support of certain result parameters, since conformance expectations for these types of parameters are implementation-specific and/or conditionally useful (e.g., `_count` parameter when an implementation expects a high volume of documents could be returned).

3.4.7 Document Search Patterns: CA:FeX-2B (Search Against Hybrid Documents Repository)

This pattern is intended to represent the more scalable method of searching for FHIR Documents from a repository that supports a hybrid model for document retrieval.

Earlier versions of this specification circulated a direct query method against a DocumentReference endpoint under this pattern. However, this has been modified in this release (in response to feedback) to better convey the use of FHIR operations as the more lightweight alternative to conventional searching against the DocumentReference endpoint (a noted challenge with the current MHD pattern).

Implementers are encouraged to consider the perspective of Data Requesters that want to have a single way for searching for both FHIR-assembled and FHIR-enabled documents. FHIR Operations present the most promise in this space but currently there are a number being explored for document exchange.

This specification has identified two promising operations (`$docRef` and `$summary`) for this pattern. We anticipate that both operations would require further enhancement to account for the current search requirements and capabilities of existing document exchange systems. For this reason, CA:FeX is seeking feedback on which of these operations Canadian implementations would be most likely to use to support FHIR Document Exchange.

`$docRef` Operation

Industry-leading implementations like the [US Core](#) and [International Patient Access \(IPA\)](#) are promoting the use of operations like `$docRef` that can generate the DocumentReference "shell" on-demand if the server contains any FHIR Documents that meet the criteria but do not already have this additional layer of metadata. Implementations that do not want to perform the secondary retrieval for FHIR Documents can choose to augment their implementation to support parameters like `_include` to return the referenced FHIR resources alongside the generated DocumentReference.

The `$docRef` operation is [currently defined in IPA](#).

Input Parameters:

- `patient.id` is currently a mandatory input parameter for the operation (see note below)
- `dateTime` is used as an optional input parameter to indicate the start and/or end dates of the care related records in the FHIR Document. If neither a start date nor an end date is provided, the most recent or current document is in scope.
- `type` is used as an optional input parameter to indicate the type of document being requested
- `on-demand` is an optional input parameter to indicate if the client wants only "on-demand" documents or both "on-demand" and "stable" documents
- `profile` is the last optional input parameter to indicate that the client only wants to retrieve documents according to a particular profile

Operation Output:

If a FHIR Document can be found that matches the provided input parameters, a Bundle with type of searchset (that includes the DocumentReference) will be returned

`$summary` Operation

Specifications like the [International Patient Summary \(IPS\)](#) have surfaced the `$summary` operation as well. This operation asks for the latest appropriate summary (whether that already exists or must be newly generated) and returns the latest summary in the form of a FHIR Document. However, this is configurable to claim other document profiles. One benefit of this approach over `$docRef` is that it circumvents the need to support an `_includes` parameter or perform a secondary retrieval step for the FHIR Document.

The `$summary` operation is [currently defined in IPS](#), but is being considered for broadening to return other FHIR Documents as well.

Input Parameters:

- `patient.identifier` is currently an optional input parameter for the operation (see note below)
- `profile` is used as an optional input parameter to indicate that the client only wants to retrieve documents according to a particular profile
- `graph` is used as an optional input parameter to indicate a URL for a GraphDefinition resource that explicitly defines the contents of the bundle returned

Operation Output:

If a FHIR Document can be found or created that matches the provided input parameters, a Bundle with type of document will be returned

Considerations for using FHIR Operations for Document Search

Both proposed operations currently rely on the Patient resource - which would necessitate that server of the Data Responder has a patient resource to anchor around. This presents a challenge to implementers who do not have plans to support a Patient resource outside of the contained content in the FHIR Document. This specification is actively seeking feedback from implementations who are considering this approach - on whether further customization is necessary.

Both operations have flexibility in whether the FHIR Document is persisted on the requester and responder servers, however the \$docRef operation is more commonly used to save a DocumentReference and Bundle on the responder server.

The \$docRef operation provides an effective option to implementers who would prefer to retrieve a list of potential documents without having to retrieve the documents in their entirety. However, use of DocumentReference is not limited to this use case and use of this operation also supports DocumentReferences that contain FHIR Binary resources in the attachment content.

3.4.8 Cutting Edge Capabilities

FHIR Operations

Following feedback from the community, Two FHIR operations that previously were included as cutting edge capabilities were brought up into the CA:FeX-2B pattern to create further differentiation between the standard MHD patterns and the lightweight patterns requested for CA:FeX.

3.5 Retrieve a Document

Scope

This capability supports the ability to retrieve known documents using the document's resource id.

3.5.1 Use Cases

The following CA:FeX use cases leverage this capability:

- [UC-02 Query and Retrieve Document](#)

3.5.2 CA:FeX Transactions

This capability describes the FHIR implementation of the CA:FeX Retrieve Data (CA:FeX-3) transaction found in the [Sequence Diagram for UC:02: Query and Retrieve Document](#).

3.5.3 HTTP Operations

This capability can be leveraged using HTTP GET requests. Implementers are expected to familiarize themselves with how searches are executed using the [FHIR RESTful API framework](#).

See [CA:FeX Actors and Transactions](#) for more information on the supported HTTP operations.

HTTP GET

In the simplest case, a retrieve is executed by performing a GET operation in the RESTful framework:

```
GET [base]/[resourcetype]/id
```

3.5.4 Considerations for Document Retrieve Patterns

This pattern is predicated on the requestor knowing the id of the resource they are attempting to retrieve.

Depending on the maturity and scope of the repository, the endpoints supported may be isolated to `[base]/Bundle` but the repository could potentially support additional document formats which would be retrieved from different endpoints.

Guidance is provided below anticipating that some adopters will solely support the ability to retrieve FHIR Documents directly, while others may choose to augment their capabilities to support a hybridized model that supports FHIR Documents alongside documents in legacy formats (e.g., Binary). This augmentation represents a different pattern that could be applied, these patterns are differentiated below using an enumerated convention (e.g., 3A, 3B).

3.5.5 Document Retrieve Patterns

CA:FeX-3A (Retrieve Document From FHIR Assembled Documents Repository)

The most straightforward method to retrieve a FHIR Document is to perform a read interaction using HTTP GET `[base]/[type]/[resource id]` against a `[base]/Bundle` endpoint.

Per the FHIR Standard on the [HTTP read interaction](#):

- This returns a single instance with the content specified for the resource type.
- This URL may be accessed by a browser.
- The possible values for the [Logical Id](#) ("id") itself are described in the [id type](#).
- The returned resource SHALL have an `id` element with a value that is the `[id]`.

Considerations for Document Retrieve Using Bundle Endpoints

Anchoring retrievals around a Bundle endpoint in this retrieval pattern is recommended for implementers that want to retrieve the FHIR Document using the known id for the Bundle. This retrieval pattern is expected to be used by implementers re-retrieving a Bundle that was previously identified in a search/or prior retrieval.

CA:FeX-3B (Retrieve Document from Hybrid Documents Repository)

Implementers are encouraged to incorporate patterns that allow for more comprehensive support for both FHIR-assembled and FHIR-enabled documents. Implementers that extend their capabilities can subsequently expect to support retrieval of documents using additional FHIR endpoints.

This pattern anchors retrieval around `[base]/DocumentReference` and/or `[base]/Binary` endpoints using the HTTP GET command.

Anchoring retrievals around a DocumentReference endpoint in this retrieval pattern is recommended for implementers that know the document they are retrieving is an embedded attachment in DocumentReference. This pattern can also be used to retrieve a DocumentReference if it was created and persisted through the use of the `$docRef` operation (see CA:FeX 2B), however this is not its primary function.

Anchoring retrievals around a Binary endpoint in this retrieval pattern is recommended for implementers that know the document that they are retrieving is a Binary resource. This retrieval pattern is expected to follow a CA:FeX Search Against Hybrid Documents Repository (CA:FeX-2B), if the `docRef` operation is used and the `attachment.url` identifies the document can be retrieved from a FHIR Binary endpoint.

This pattern is anticipated to be further fleshed out as use cases requiring this pattern arise.

3.5.6 Document Retrieve Patterns: Supported Retrieval Parameters

The following examples are used to retrieve documents based on a known resource ID.

GET `[base]/Bundle/[id]`

GET `[base]/DocumentReference/[id]`

GET `[base]/Binary/[id]`

3.5.7 Cutting Edge Capabilities

FHIR Operations

[FHIR Operations](#) are used (a) where the server needs to play an active role in formulating the content of the response, not merely return existing information, or (b) where the intended purpose is to cause side effects such as the modification of existing resources, or creation of new resources.

This specification does not currently define any custom operations for the Retrieve for a Document Capability. Unlike [Submit a Document](#) and [Search for a Document](#) capabilities, there are no existing operations that this specification has identified that directly impact retrieval of a document resources using resource id.

3.6 Response Handling

3.6.1 HTTP Status Codes

The HTTP status codes returned in the response aid applications in understanding whether a search was successful, and if not successful, provides some context as to why the search did not return the expected response.

These status codes are considered the base status codes that should be present in any implementing API. Implementations and guides may include additional status codes that their service supports and may supply further details on the circumstances or expected client behavior.

Error

If the search fails (cannot be executed, not that there are no matches), the return value SHALL be a status code 4xx (client error) or 5xx (API/service error). An `OperationOutcome` SHOULD be returned detailing the error.

| Scenario | HTTP Status Code | Outcome |
|--|--------------------------------|--|
| The request is performed with incorrect syntax or against a resource or search parameter that the API does not currently support. | HTTP 400 Bad Request | Error code is returned in the response with an <code>OperationOutcome</code> detailing the error |
| The request is performed by an application that has not properly authenticated. | HTTP 401 Unauthorized | Error code is returned in the response with an <code>OperationOutcome</code> detailing the error |
| The request is performed by an authenticated user, but the authenticated user is not permitted to perform the requested operation. | HTTP 403 Forbidden | Error code is returned in the response with an <code>OperationOutcome</code> detailing the error |
| The request is performed on a resource type that is not supported | HTTP 404 Not Found | Error code is returned in the response with an <code>OperationOutcome</code> detailing the error |
| The request is performed using an HTTP method that the server does not support for the resource. | HTTP 405 Method Not Allowed | Error code is returned in the response with an <code>OperationOutcome</code> detailing the error |
| The request includes a media type which is not supported. For example, the client uploads an image as <code>image/svg+xml</code> , but the server requires that images use a different format. | 415 Unsupported Media Type | Error code is returned in the response with an <code>OperationOutcome</code> detailing the error |

| Scenario | HTTP Status Code | Outcome |
|--|---------------------------|---|
| The request is performed but the server encountered an internal error during the process of the response message. | 500 Internal Server Error | Error code is returned in the response with an OperationOutcome detailing the error |
| The request is made when the service is temporarily unavailable. | 503 Service Unavailable | Error code is returned in the response with an OperationOutcome detailing the error |
| The request is performed but the server, while acting as a gateway or proxy, does not get a response in time from an upstream server in order to complete the request. | 504 Gateway Timeout | Error code is returned in the response with an OperationOutcome detailing the error |

Success

If the search is properly formatted and succeeds, the API SHALL return an HTTP 200 OK status code and the return content SHALL be a Bundle with type = searchset containing the results of the search as a collection of zero* or more resources in a defined order.

NOTE: Searches that are properly performed but have no PHI to return will still result in a HTTP 200 OK status code

| Scenario | HTTP Status Code | Outcome |
|--|------------------|---|
| The request is performed with correct syntax and there are FHIR resources available that match those search parameters | HTTP 200 OK | Success code is returned in the response, Bundle returned in the response contains entries with requested FHIR resources |
| The request is successful and a FHIR resource has been created on the server | HTTP 201 Created | Success code is returned in the response, the location of the resource is returned in the location header of the response |

3.6.2 Operation Outcome

OperationOutcome is used to provide more detailed description of any issues that occurred during execution of an operation. The general conventions for OperationOutcome are:

- OperationOutcome can either be the sole response to an operation (generally accompanying an HTTP code that indicates a failure) or it can be part of a Bundle indicating potential warnings associated with the generation of the search response.
- If an OperationOutcome is returned with anything other than a success (200), the issues it contains will be of type 'error' or 'fatal'. Fatal is used for issues that occur before the query can be exercised, 'error' is used for issues that occur during query execution.

- OperationOutcomes returned as part of a Bundle will only contain 'warning' and 'information' messages. The user should always be made aware of the existence of these messages and the user should have the opportunity to review all such messages.
- The OperationOutcome.issue.code provides a standardized description of the issue. Systems MAY create logic based on the specific code returned.
- The details of the issue or warning will be found in issue.details.text. This content should always be displayed to the user.
- The issue.details.coding values, issue.diagnostics and issue.location are intended for diagnostic purposes and will generally only be relevant to help-desk personnel. It may be appropriate to make access to this information available only by clicking on a button rather than displaying information to the user that may be confusing.
- issue.location will only occur if the problem is with a submitted FHIR instance (i.e. it will not be present if the issue is for query parameters, HTTP headers, etc.). It will be expressed as an XPath regardless of whether the submitted content is XML or JSON.

4 Reference Architecture

4.1 Purpose

This Reference Architecture provides guidance on how to apply CA:FeX implementation patterns along with other IHE profiles to address interoperability needs pertaining to FHIR Health Information Exchange (HIE) implementation pattern. It defines the rules of engagement and defines the interoperability landscape as well as developing a connectivity platform to allow external vendors to test solutions.

4.2 Intended Audience

The intended audience of this Reference Architecture includes, but is not limited, to the following:

- IT departments of healthcare institutions (technical product managers, IT managers, operations staff)
- Technical staff of clinical solution vendors
- Experts involved in standards development
- Individuals and teams responsible for implementing software solutions such as project managers, CTOs, CISOs, software engineers, technical product managers, IT managers, operations staff, and other similar roles.

4.3 Overview

This section contains a list of recommended IHE candidate profiles that can potentially be grouped with CA:FeX that can meet specific needs for FHIR Health Information Exchange (HIE) implementation pattern. The sequence diagrams included in this section group together actors and transactions from multiple profiles including CA:FeX to address the scope of the use cases.

4.4 How To Read The Reference Architecture

- This document lists actors and transactions from candidate IHE profiles that can potentially be grouped with the CA:FeX Interoperability Specification and enable FHIR Health Information Exchange (HIE) implementation patterns.
- The relationship between consuming and recipient actors are indicated by black lines.
- There are two swim lanes that group together the actors and transactions: Clinical systems and Jurisdictional systems. It is assumed that vendors will take on roles from the clinical systems while provinces and territories will take on roles from the jurisdictional systems.
 - Dark green actors are required actors showcasing different roles within the clinical systems that are required for this release
 - Light green actors are optional actors showcasing different roles within the clinical systems that are optional for this release
 - Dark blue actors are required actors showcasing different roles within the jurisdictional systems that are required for this release
 - Light blue actors are optional actors showcasing different roles within the jurisdictional systems that are optional for this release
- The black lines are the interactions and standardized transactions of the actors and their association with one another as defined by the IHE methodology framework.
- A list of all the IHE profiles are included at the bottom of the Reference Architecture.
- A list of pan-Canadian Profiles are listed at the bottom of the Reference Architecture.

- A legend is found at the bottom of the Reference Architecture to help readers orient themselves to the diagram.

4.5 How To Use The Reference Architecture

The list below summarizes how to use this document:

- **Role Identification:** Jurisdictions and vendors will need to identify their role (e.g. actors) from the Reference Architecture and sequence diagrams for each of the use cases in scope.
- **Gap Identification:** Based on the role(s) identified from the Reference Architecture and sequence diagrams, potential assessment is needed for identification of gaps for meeting the requirements of the standardized actors and transactions needed to satisfy particular use cases.
- **Provincial Reference Architecture:** Provinces and jurisdictions may need to draft their own version of Reference Architecture specific to their needs. Current technology landscape, existing architecture and current business priorities will help in developing a version for the province.
- **Document Evolution and Feedback:** This is a living document and will evolve based on feedback and refinements to the the CA:FeX Interoperability Specifications and related use cases. This document is published on InfoScribe to capture comments and feedback from all key stakeholders. Additionally, multiple sessions will be conducted to discuss and update the content of this document.
- **Vendor Conformance Testing (Connectathon / Projectathon):** This document will provide an opportunity for vendors to prepare for conformance testing of the CA:FeX Interoperability Specifications on the IHE Gazelle platform. IHE Gazelle is an open-source, web-based test platform supporting a wide portfolio of interoperability test tools suited to validate interface conformity to IHE Profiles and project-specific standards-based interoperability specifications. Vendors can validate their products and eHealth projects to procure interfaces they deploy. For additional information on Gazelle, please refer to the following link: [IHE Gazelle](#)

*Note: It's expected that the reader should have a moderate degree of familiarity with IHE profiles, esp. ATNA, CT and IUA.

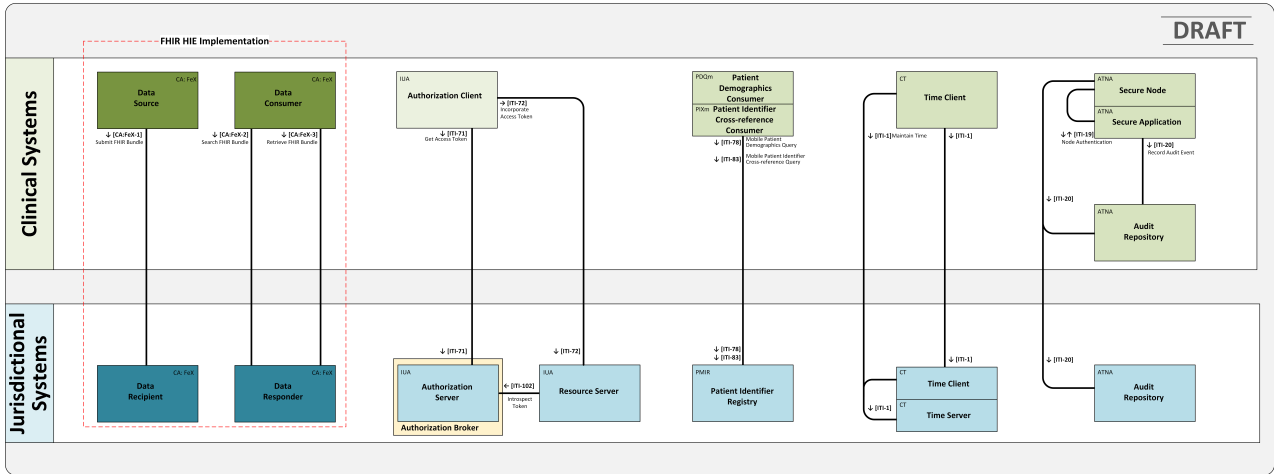
4.6 High Level View - Reference Architecture for CA:FeX

This is a high-level view of relevant IHE Integration Profiles to support the FHIR HIE implementation using CA:FeX. Required and optional capability support is described in the sequence diagrams associated with each Use Case analysis. Vendors or jurisdictions may decide not to implement optional IHE profiles showcased below, however it is highly recommended that areas pertaining to authentication, auditing and security are being addressed using solutions that currently exist in their respective enterprise architecture.

CA:FeX Health Information Exchange




DRAFT



- IHE Profiles**
- IUA Internet User Authorization
 - PIXm Patient Master Identity Registry
 - PODm Patient Demographics Query for Mobile
 - PIXm Patient Identifier Cross-Reference for Mobile
 - XDM Cross-enterprise Document Media Interchange
 - ATNA Audit Trail and Node Authentication
 - CT Consistent Time

- pan-Canadian Interoperability Specification Profile**
- CA:FeX
 - FHIR Exchange



4.7 Foundational IHE Profiles

Background

Foundational IHE Profiles address critical interoperability issues such as user authorization (e.g., IUA), security node and audit records (e.g., ATNA) and consistent time (e.g., CT).

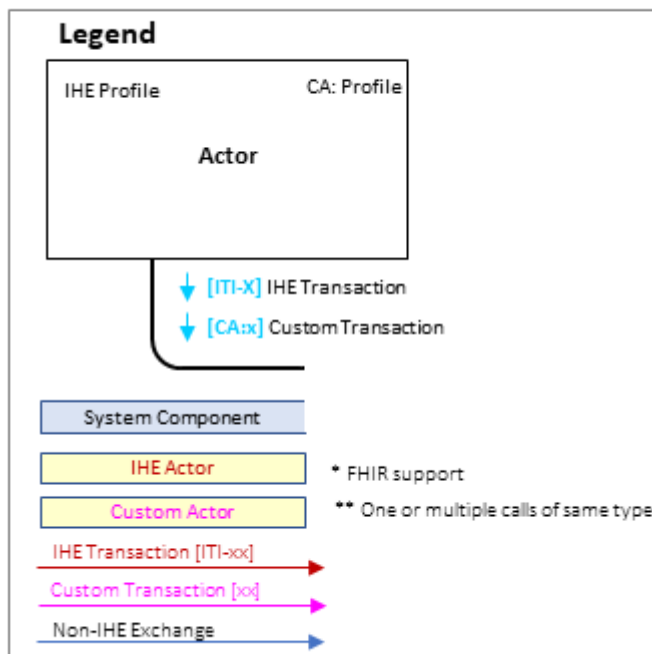
Assumption

Vendors and jurisdictions in the ecosystem can optionally choose to play the standardized actors and transactions listed in the Foundational Profiles to support implementation of the CA:FeX and to address authentication, auditing and security needs. Additional information and requirements for these Foundational Profiles can be found below. Vendors or jurisdictions may decide not to implement optional IHE profiles listed below, however it is highly recommended that areas pertaining to authentication, auditing and security are addressed using solutions that currently exist in their respective enterprise architecture.

IHE Profiles Profiles included: IUA, ATNA, CT

Legend

The following diagram is the legend for the sequence diagrams to help readers orient themselves with the diagrams.

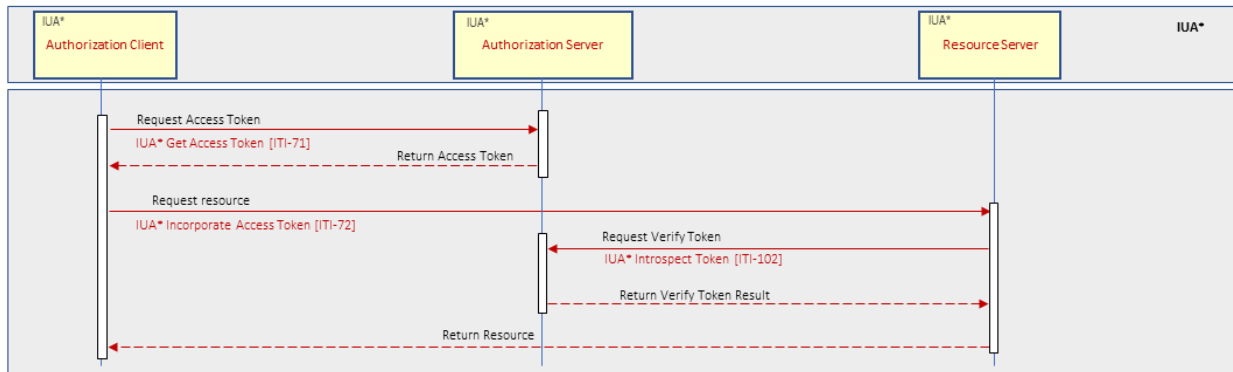


IUA*

The IUA* (Internet User Authorization) provides support for authorization to access resources when using HTTP RESTful transports, by managing authorization tokens.

The Authorization Client must have a valid token that must be presented to resource server with every request:

- Get Access Token [ITI-71] – performed when Client does not yet have a token or when token is expired
- Incorporate Access Token [ITI-72] – the client must include the token with every request
- Introspect Token [ITI-102] – the resource server must introspect token at every request



ATNA*

ATNA* (Audit Trail and Node Authentication) provides support for ensuring that the communicating systems have a level of trust in each other through node authentication, that communications between the different system components are encrypted (via TLS), and that system activity is audited.

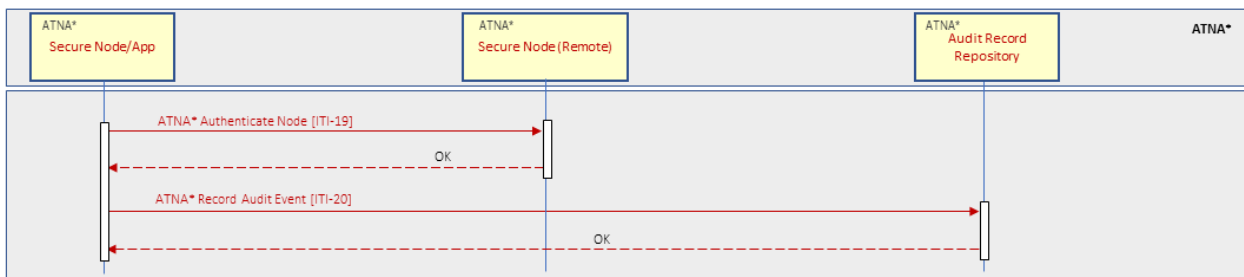
- ATNA* Authenticate Node [ITI-19]

Before establishing secure communication, mutual authentication is performed between two secure nodes. A secure pipe will be established through which secure transactions will take place.

Secure Node also authenticates the identity of the user who requests access to the node.


- ATNA* Record Audit Event [ITI-20]

The Secure Node/App sends auditable events to an Audit Record Repository. The triggers for sending audit logs can vary and may be specified in other IHE profiles, local law or regulation, or local policy.



CT

CT (Consistent Time) ensures that the system clocks and time stamps of the many computers in a network are well synchronized. Synchronization with a median error less than 1 second is sufficient for most purposes.



4.8 Sequence Diagram for UC-01: Create and Submit Document

Scenario: Clinical Solution A Retrieves Document from a Central Clinical Data Repository

Assumption: Document is stored in a Clinical Data Repository

This sequence diagram provides the option of using the CA:FeX Interoperability Specifications that provide support for saving and retrieving data (e.g., documents) to and from a central Clinical Data Repository. This specification includes a Data Source and a Data Recipient actor. Additionally, this sequence diagram uses the 'Submit Data' FHIR operation.

Sequence Diagram Overview:

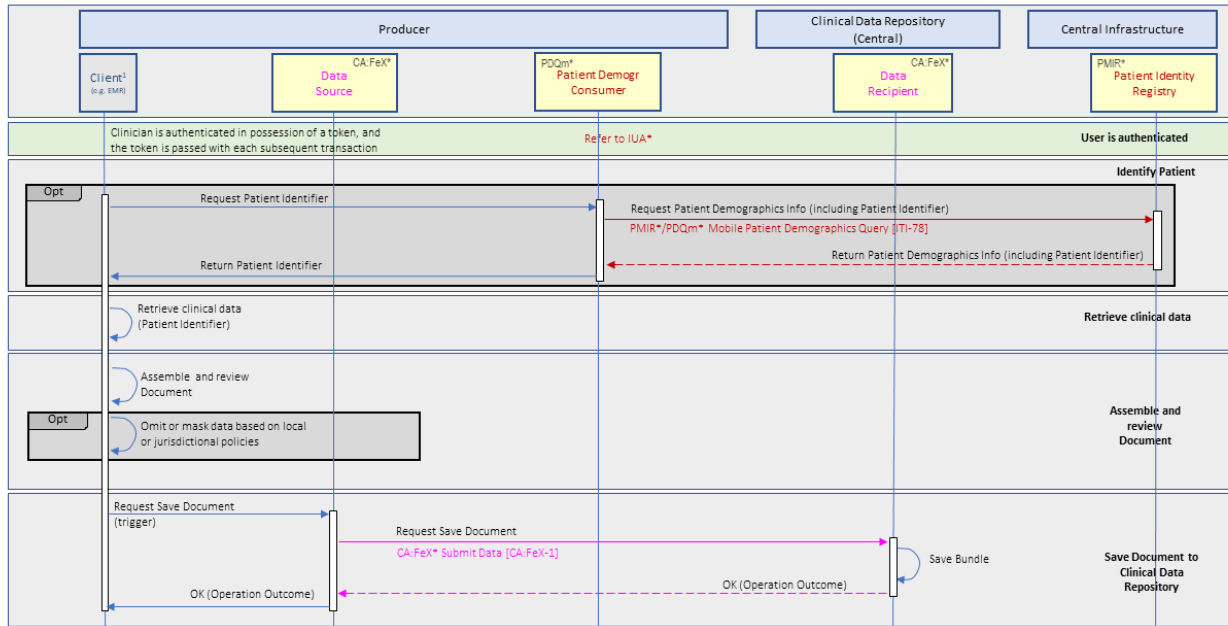
Below provides guidance on how to read the sequence diagram:

- This sequence diagram illustrates how the different standardized actors of a system should interact with each other to carry out specific standardized transactions, and the order in which the transactions and interactions occur when UC-01 of the CA:FeX Specification is executed.
- The legend on the bottom right corner describes the different system components, actors and transactions that are necessary to carry out this use case.
- The green swim lane is a simplified view of the actors and transactions required by the Foundational IHE Profiles, in addition to the other ones that are not explicitly illustrated on the diagram (e.g., ATNA, CT, etc.) but included as a white note. These are pre-requisite conditions for this use case and it is assumed that these will be satisfied.
- The blue swim lanes groups a sequence of processes (along with their required actors and transactions) that are needed to occur to satisfy this use case. These are to be read from left to right and top to bottom.
- The red note boxes describe important information and notes that provide more context for the sequence diagram.
- More information about the details of the Foundational IHE Profiles can be found in the PS-CA Specifications available [here](#).

Additional Considerations

The sequence diagrams included in this section do not showcase all of the possible combinations of IHE profiles and transactions that can be used for a particular implementation pattern. For example, ITI-83 transaction can be used in place of ITI-78 if the preferred implementation pattern is PIXm/PMIR.

UC-01: Create and Submit Document – CA:FeX

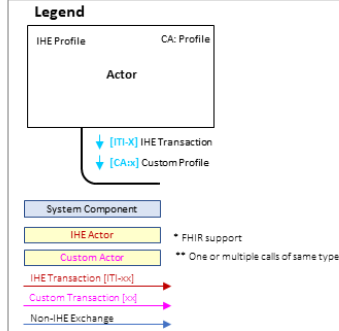


Footnotes

- Examples of the client system can be any of the following: EMR, HIS, CIS, PHR, or EHR

Prerequisites

- Client is logged into the system ([UA*])
- Client obtained a valid access token from the Authorization Server that is used with each transaction([UA*])
- All communication is done through secure channels (ATNA)
- System time is synchronized among all components (CT)



4.9 Sequence Diagram for UC:02: Query and Retrieve Document

Scenario: Clinical Solution A Retrieves Clinical Data from a Clinical Data Repository.

Assumption: Clinical Data is stored in a Central Clinical Data Repository

This sequence diagram provides the option of using the CA:FeX Interoperability Specifications that provide support for saving and retrieving data (e.g. documents) to and from a central Clinical Data Repository. This specification includes a Data Consumer and a Data Responder actor. Additionally, this sequence diagram uses the 'Search Data' and 'Retrieve Data' FHIR operations.

Sequence Diagram Overview:

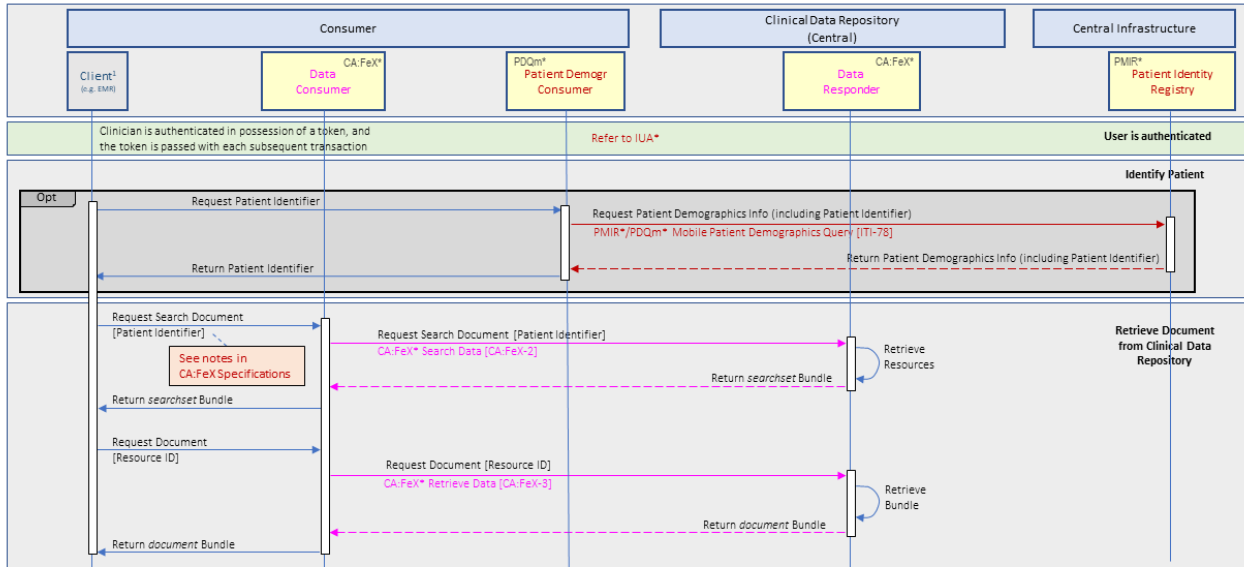
Below provides guidance on how to read the sequence diagram:

- This sequence diagram illustrates how the different standardized actors of system should interact with each other to carry out specific standardized transactions, and the order in which the transactions and interactions occur when Use Case 2 of the CA:FeX Specification is executed.
- Data Consumer role varies by client type:
 - HCP: Clinical Solution, e.g., EMR used by the Health Care Provider to request & retrieve access to clinical document(s) from the Clinical Data Repository
 - Patient: Patient Portal used by the Patient / Subject of Care to request & retrieve access to their clinical document(s) from the Clinical Data Repository
- The legend on the bottom right corner describes the different system components, actors and transactions that are necessary to carry out this particular use case.
- The green swim lane is a simplified view of the actors and transactions required by the Foundational IHE Profiles, in addition to the other ones that are not explicitly illustrated on the diagram (e.g. ATNA, CT) but included as a white note. These are pre-requisite conditions for this particular use case and it is assumed that these will be satisfied.
- The blue swim lanes groups sequence of processes (along with their required actors and transactions) that are needed to occur to satisfy this particular use case. These are to be read from left to right and top to bottom.
- The red note boxes describe important call outs, information and notes that provide more context for the sequence diagram.
- More information about those details of the Foundational IHE Profiles can be found in the PS-CA Specifications available [here](#).

4.9.1 Additional Considerations

The sequence diagrams included in this section do not showcase all of the possible combinations of IHE profiles and transactions that can be used for a particular implementation pattern. For example, ITI-83 transaction can be used in place of ITI-78 if the preferred implementation pattern is PIXm/PMIR.

UC-02 Query and Retrieve Document – CA:FeX



Footnotes

1. Examples of the client system can be any of the following: EMR, HIS, CIS, PHR, or EHR

Prerequisites

- Client is logged into the system (UA*)
- Client obtained a valid access token from the Authorization Server that is used with each transaction (UA*)
- All communication is done through secure channels (ATNA)
- System time is synchronized among all components (CT)

